

Performance Improvement of LIBRSB-1.3 Sparse BLAS on EuroHPC-class Systems

Michele MARTONE* (michele.martone@lrz.de)

Leibniz Supercomputing Centre, Garching bei München, Germany

Context: Project LyNcs

“project “LyNcs”: “Linear Algebra, Krylov-subspace methods, and multi-grid solvers for the discovery of New Physics”

subproject of PRACE-6IP Grant agreement ID: 823767, Work Package 8 (“Forward-looking software solutions towards Exascale”)

synergy of expertises and codes spanning the entire software stack:

- LQCD, The Cyprus Institute: DDALPHAAMG, LYNCS-API
- numerical linear algebra, Inria Bordeaux (France): FABULOUS
- portable performance kernels, LRZ: LIBRSB

this poster: main LIBRSB achievements during the project

Numerical Techniques of Interest to LyNcs

- iterative methods: *block Krylov*
- require efficient Sparse Matrix-Matrix multiplication aka SpMM

SpMM in matrix form (m aka *NRHS* aka *number of right hand sides*):

$$\begin{bmatrix} c_{11} & \dots & c_{1m} \\ \vdots & \ddots & \vdots \\ c_{n1} & \dots & c_{nm} \end{bmatrix} \leftarrow \beta \begin{bmatrix} c_{11} & \dots & c_{1m} \\ \vdots & \ddots & \vdots \\ c_{n1} & \dots & c_{nm} \end{bmatrix} + \alpha \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nn} \end{bmatrix} \begin{bmatrix} b_{11} & \dots & b_{1m} \\ \vdots & \ddots & \vdots \\ b_{n1} & \dots & b_{nm} \end{bmatrix}$$

More on the methods of our interest:

R. B. Morgan. *Restarted block GMRES with deflation of eigenvalues*, Appl. Numer. Math., 54(2):222–236, 2005. • M. Robbe and M. Sadkane. *Exact and inexact breakdowns in the block GMRES method*, Linear Algebra Appl., 419:265–285, 2006. • E. Agullo, L. Giraud, and Y.-F. Jing. *Block GMRES method with inexact breakdowns and deflated restarting*, SIAM J. Matrix Anal. Appl., 35(4):1625–1651, 2014.

Recursive Sparse Blocks (RSB) Layout

- for large matrices (uses *cache locality*, *coarse thread parallelism*)
- supports *autotuning* (layout adjusted to maximize performance)




Figure 4: Instance of classical test matrix *bayer02* ($14k \times 14k$, 64k nonzeros). Black-bordered boxes are *sparse blocks*, and are *Z-ordered*. Greener have fewer nnz than average, redder have more. Either in “Coordinate” format (COO) or “Compressed Sparse Rows” (CSR). Blocks rows (LHS) and columns (RHS) ranges evidenced (left and top side).

Setup of SpMM Performance Experiment

- Five machines with different architectures
 - amd64: smng (Intel Skylake), rome (AMD Rome), coop (Intel Cooper Lake)
 - arm64: thx (Marvell ThunderX2), a64fx (Fujitsu A64FX)
- Same input data and parameter ranges
 - 24 threads, autotuning on
 - 1,2,4 NRHS (*right hand sides* operands to SpMM)
 - Operands layouts: by-rows and by-columns
 - The four BLAS numerical types
 - A selection of 44 matrices (symmetric and unsymmetric)
- Compilation flags
 - smng,rome: `icc -ipo -O3 -no-prec-div -fp-model fast=2 -xHost`
 - gcc -Ofast -march=native -mtune=native elsewhere

Conclusions

- The LyNcs project allowed improving LIBRSB in:
 - Performance of SpMM (*by-rows* is optimal operands layout now)
 - Usability via C++, GNU OCTAVE, PYTHON
 - Many aspects of robustness and overall quality (e.g. tests, coverage, continuous integration — details out of the scope of this poster)

LIBRSB

project page: <http://librsb.sf.net>

- >100KLOC of C99, OPENMP, and modern templated C++
- node-level **shared-memory**-parallel operations, for:
 - Sparse BLAS: matrix assembly/destroy, SpMM, triangular solve
 - interactive applications (update, sparse-sparse ops, conversions)
 - distributed-memory applications (block extract, update, etc.)
- with own interface in C/C++ and FORTRAN
 - with Sparse BLAS interface (BLAS Technical Forum Standard)
 - with interfaces for GNU OCTAVE and PYTHON interpreters
- LGPLv3-licensed free software, available via SPACK, GUIX-HPC, EasyBuild, and on Debian, Ubuntu, OpenSUSE, Windows,...

New in LIBRSB-1.3: Modern C++ interface

```
1 #include <rsb.hpp>
2 #include <vector>
3 #include <array>
4 using namespace ::rsb;
5
6 int main() {
7     RsbLib rsb;
8     const int nnzA { 7 }, nrhs { 2 };
9     const int nrA { 6 }, ncA { 6 };
10    const std::vector<int> IA {0,1,2,3,4,5,1};
11    const int JA [] = {0,1,2,3,4,5,0};
12    const std::vector<double> VA {1,1,1,1,1,1,2}, B(nrhs * ncA,1);
13    std::array<double,nrhs * nrA> C;
14    const double alpha {2}, beta {1};
15
16    RsbMatrix<double> mtx(IA,JA,VA,nnzA); // Notice above declarations of IA,JA,VA
17
18    mtx.spmv(RSB_TRANSPOSITION_N, alpha, nrhs, RSB_FLAG_WANT_ROW_MAJOR_ORDER, B,
19            beta, C); // ditto for B and C
20 }
```

Figure 1: Program using the new modern C++ interface to create a matrix and multiply it by a vector. Usage of C++20’s `std::span` in the declaration of `RsbMatrix()` and `spmv()` ensures type safety while not prescribing usage of any specific array type.

Improved in LIBRSB-1.3: SpMM + Fortran Layout





Figure 5: A matrix and its two SpMM operands stored in **column-major** (also known as FORTRAN) order. The matrix consists of four *sparse blocks*, of which one is highlighted. Each operand is stored one column after the other, with memory following the blue line. Consequently, each operand’s two column portions pertaining a given sparse block are **not contiguous**. LIBRSB-1.2 would multiply each sparse block by each contiguous subvector. LIBRSB-1.3 can now perform this in one pass only.

GNU Octave + liboctave + LIBRSB = SparseRSB



- GNU OCTAVE: a MATLAB-like interactive numerical language
- liboctave: access OCTAVE via C++

```
1 octave:1> A = spsersb ( sparse ( rand (3) > .6) )
2 A =
3
4 Recursive Sparse Blocks (rows = 3, cols = 3, nnz = 2 [22%])
5
6 (3, 1) -> 1
7 (3, 3) -> 1
```

Figure 2: The `spsersb` keyword allows transparent usage of LIBRSB in the same way as the `sparse` built-in. Part of Octave Forge.

Python + Cython + LIBRSB = PyRSB



- SCI-PY: popular PYTHON scientific computing API
- CYTHON: *optimising static compiler* for C extensions to PYTHON

```
1 import numpy
2 import scipy
3 from rsb import import rsb_matrix
4 V=[ 11.,12.,22.]
5 I=[ 0, 0, 1]
6 J=[ 0, 1, 1]
7 a = rsb_matrix((V, (I, J)), [3,3])
8 ...
9 y = y + a * x;
```

Figure 3: PYRSB’s usage is styled after SciPy’s `csr_matrix`.

Improved in LIBRSB-1.3: SpMM + C Layout




Figure 6: A matrix and its two SpMM operands stored in **row-major** (also known as C) order. The matrix consists of four *sparse blocks*, of which one is highlighted. Each operand stores one row after the other, with memory following the blue line. Consequently, the two column portions operands pertaining a given sparse block are **contiguous**. Here too LIBRSB-1.2 kernels would multiply each sparse block repeatedly. LIBRSB-1.3 can perform this in one pass only. Note that for NRHS=1, by-rows and by-columns are equivalent.

Speedup of by-rows Layout over by-columns Layout

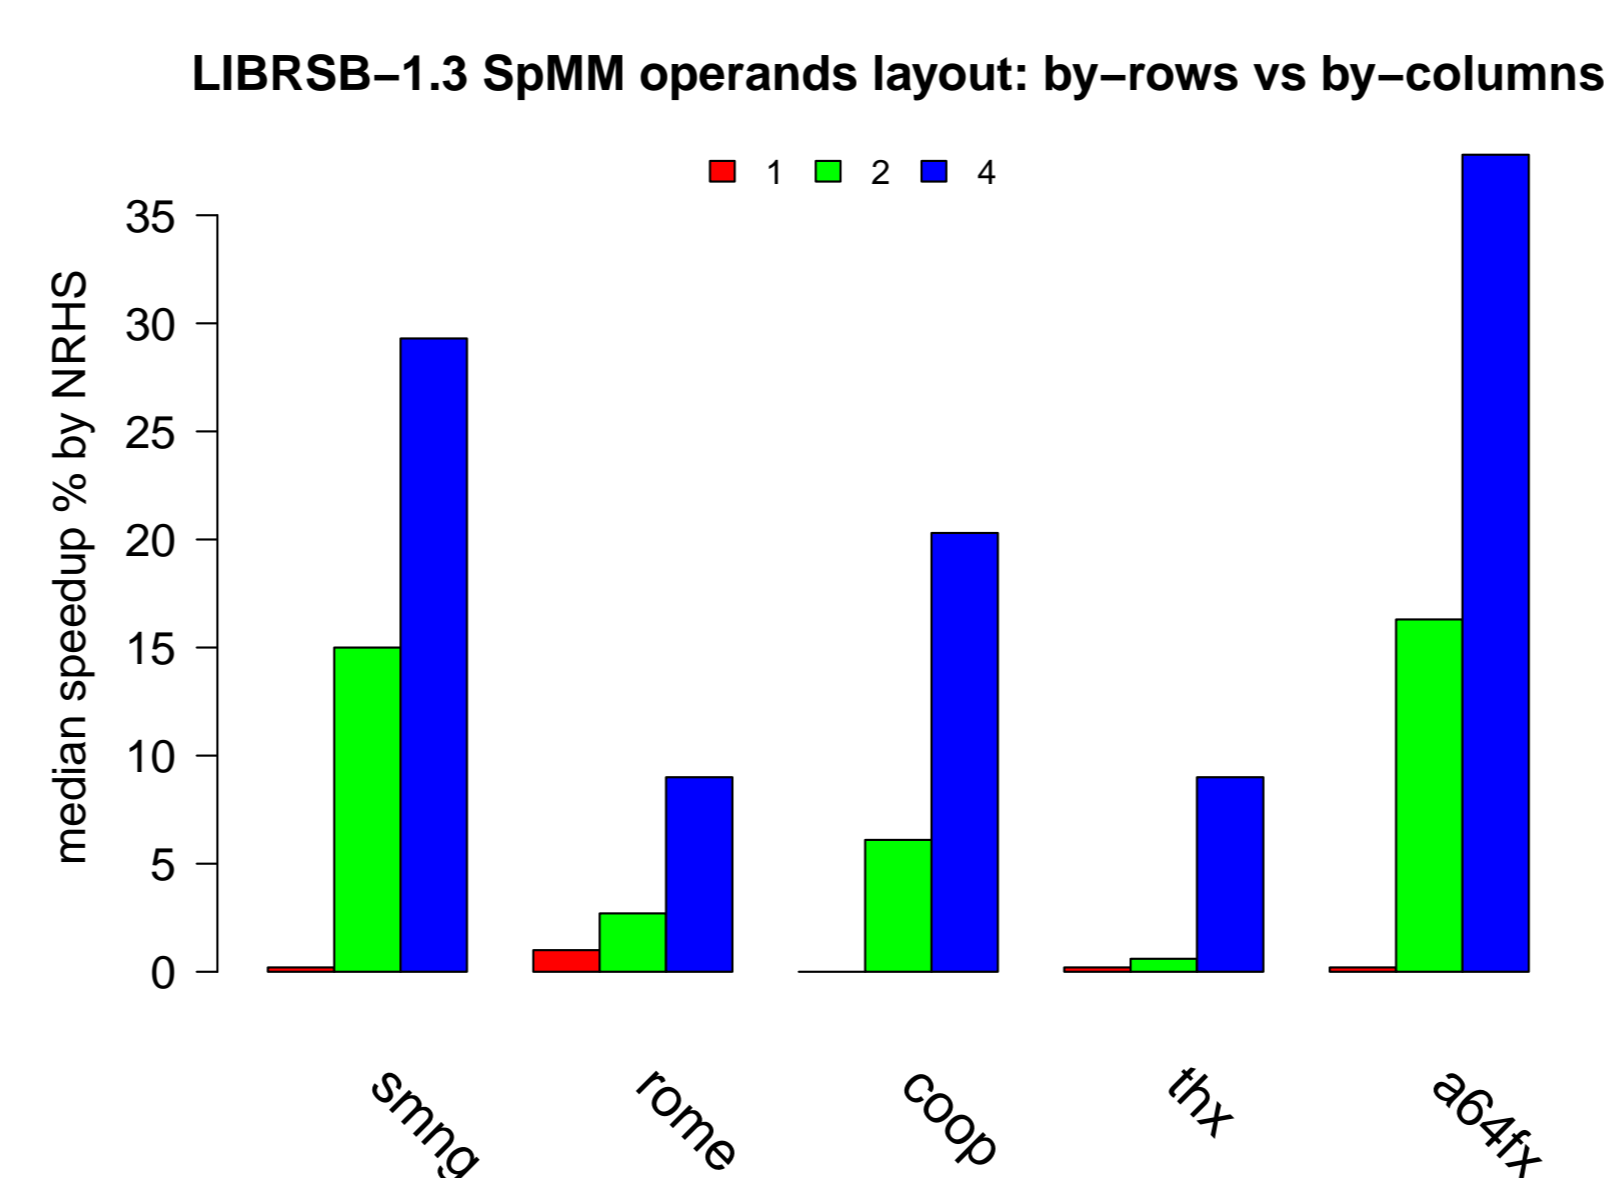


Figure 7: For each combination of matrix and numerical type, by-rows-layout SpMM performance is divided by by-columns-layout performance, and here grouped by machine and NRHS. For SpMM (NRHS>1), by-rows layout is now the recommended layout (in LIBRSB-1.2, by-columns was the optimal one). With NRHS=2 it improves up to 16%; with NRHS=4, by 9–38%, depending on the architecture.

Speedup of version 1.3 over version 1.2

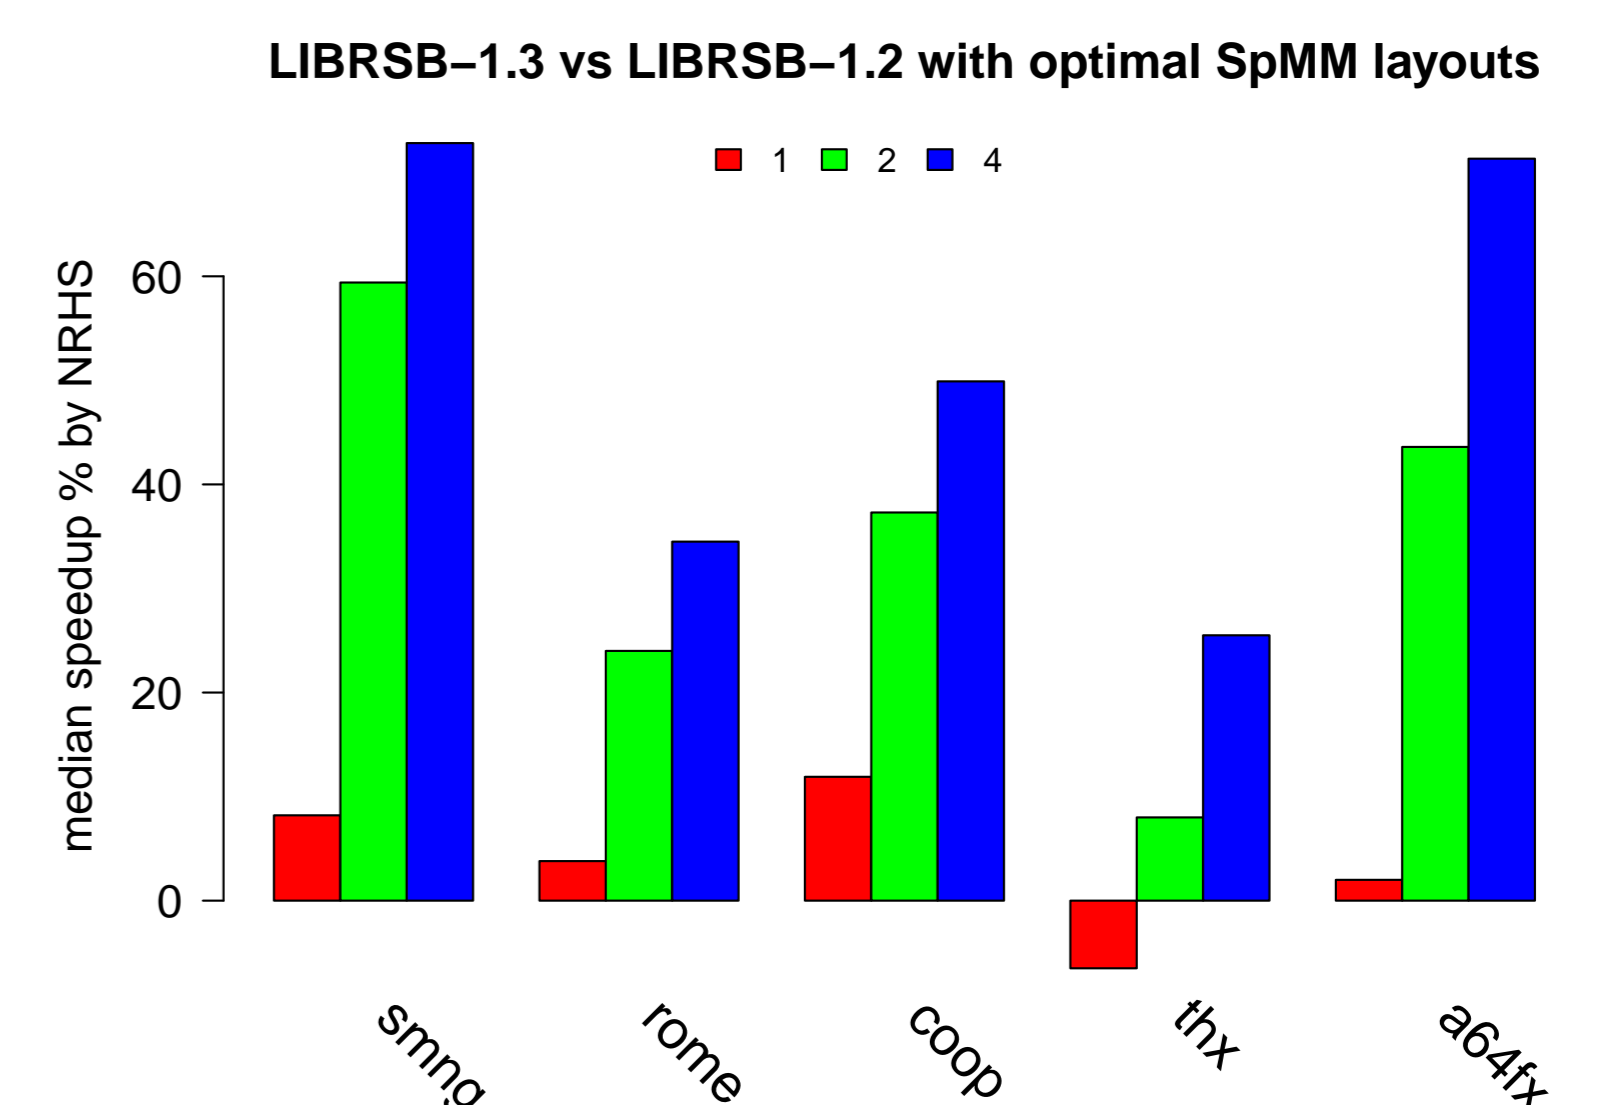


Figure 8: With the exception of SpMV on thx, 1.3-to-1.2 ratios median values indicate an SpMM improvement across all architectures. The SpMV (NRHS=1) speedup median improvement amounts to a few percentage points (for improvements obviously unrelated to the operands layout). If using the per-version recommended operands’ layout, SpMM with NRHS=2 improves by 8–59%, SpMM with NRHS=4 by 25–73%, depending on the architecture.

Acknowledgements

The Implementation Phase of PRACE receives funding from the EU’s Horizon 2020 Research and Innovation Programme (2014–2020) under grant agreement 823767.



The author gratefully acknowledges the Gauss Centre for Supercomputing e.V. (<https://www.gauss-centre.eu>) for funding this project by providing computing time on the GCS Supercomputer SuperMUC-NG at the Leibniz Supercomputing Centre (<https://www.lrz.de>). Part of the performance results have been obtained on systems in the test environment BEAST (BAVARIAN ENERGY ARCHITECTURE & SOFTWARE TESTBED) also at LRZ.

