Towards Fast, Accurate and Reproducible LU Factorization

Roman Iakymchuk, David Defour and Stef Graillat

CST/PDC, CSC, KTH Royal Institute of Technology
Lindstedtsvägen 5, 11428 Stockholm, Sweden
riakymch@kth.se

DALI–LIRMM, Université de Perpignan
52 avenue Paul Alduy, F-66860 Perpignan, France
david.defour@univ-perp.fr

Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6
4 place Jussieu, F-75005 Paris, France
stef.graillat@lip6.fr

Keywords: LU factorization, triangular solver, matrix-vector product, reproducibility, accuracy, long accumulator, error-free transformation.

The process of finding the solution of a linear system of equations is often the core of many scientific applications. Usually, this process relies upon the LU factorization, which is also the most compute-intensive part of it. Although current implementations of the LU factorization may reach high performance, their reproducibility[1] and, even more, accuracy cannot be guaranteed, mainly, due to the non-associativity of floating-point operations, the concurrent thread-level execution of independent operations on CPUs or warp scheduler on GPUs.

In this work, we address the problem of reproducibility of the LU factorization due to cancelations and rounding errors, resulting from floating-point arithmetic. Instead of developing an implementation of the LU factorization from scratch, we benefit from the hierarchical structure of linear algebra libraries and start by developing/enhancing

[1]We define reproducibility as an ability to obtain a bit-wise identical floating-point result from multiple runs of the code on the same input data.

reproducible algorithmic variants for the kernel operations – like the ones included in the BLAS library – that serve as building blocks in the LU factorization. In addition, we aim to improve the accuracy of these underlying BLAS routines.

We consider an unblocked algorithmic variant of the LU factorization that can be expressed in terms of BLAS level 1 and 2 routines: triangular solver (TRSV), dot product (DOT) and matrix-vector product (GEMV). We prevent cancellations and rounding errors in these kernels by applying a long accumulator and error-free transformations [1]. Additionally, we tackle the problem of accuracy in the reproducible triangular solver (ExTRSV) [2] through iterative refinement. Moreover, we provide an accurate and reproducible algorithmic variant as well as an implementation of the matrix-vector product. Thus, following the bottom-up approach, we construct the reproducible algorithmic variant of the LU factorization. Up to our knowledge, this is the first implementation of a reproducible LU factorization. In global, we present initial evidence that reproducible higher-level linear algebra operations can be constructed following the hierarchical bottom-up approach.

References:

[1] S. Collange, D. Defour, S. Graillat, R. Iakymchuk, Numerical Reproducibility for the Parallel Reduction on Multi- and Many-Core Architectures, *Parallel Computing*, 49 (2015), pp. 83–97.

[2] R. Iakymchuk, S. Collange, D. Defour, S. Graillat, *Reproducible Triangular Solvers for High-Performance Computing*, In the Proceedings of the 12th International Conference on Information Technology: New Generations (ITNG 2015), Special track on: Wavelets and Validated Numerics, April 13-15, 2015, Las Vegas, Nevada, USA, pp. 353-358.