DEVELOPMENT OF A LOW-LEVEL, ALGEBRA-BASED LIBRARY TO PROVIDE PLATFORM PORTABILITY ON HYBRID SUPERCOMPUTERS

Xavier Álvarez-Farré[*,1], Àdel Alsalti-Baldellou[1,2], Guillem Colomer[1], Andrey Gorobets[3], Assensi Oliva[1] and F. Xavier Trias[1]

[1] Heat and Mass Transfer Technological Center, Technical University of Catalonia, Carrer Colom 11, 08222 Terrassa (Barcelona), Spain.
[2] TermoFluids S.L. (https://www.termofluids.com), Sabadell (Barcelona), Spain.
[3] Keldysh Institute of Applied Mathematics, Russian Academy of Sciences, Miusskaya Sq. 4, 125047 Moscow, Russia.

Modern HPC systems consist of multiple hybrid computing nodes interconnected via a communication infrastructure. Hybrid nodes combine computing devices of different architectures, such as manycore CPUs or massively parallel GPUs, among others. To take advantage of such systems, the computing subroutines that form the algorithms, the so-called kernels, must be compatible with distributed- and shared-memory multiple instruction, multiple data parallelism, and more importantly, with stream processing, which is a very restrictive parallel programming paradigm. Therefore, a complex hierarchical parallel implementation model is required for combining the different parallel paradigms and the corresponding computing frameworks. This programming complexity encourages the demand for portable and sustainable implementations of scientific simulation codes. In this line, we proposed in [1] an algebra-based framework for the scale-resolution of incompressible turbulent flows on hybrid supercomputers. Briefly, the CFD algorithm relies on a set of only three algebraic kernels. Thus, the kernel code shrinks to hundreds of lines; portability becomes natural, and maintaining the OpenMP, OpenCL, and CUDA implementations requires little effort. Besides, we can easily use standard libraries optimized for a particular architecture, in addition to our specialized in-house implementations.

On the development of the aforementioned framework, we have restructured the code to implement a standalone, low-level algebraic library that can be coupled beneath any other high-level, algebra-oriented framework. Namely, this library provides platform portability in the simplest possible manner. Internally, algebraic objects are distributed among computing devices using a multilevel decomposition approach. Data exchanges between computing units or between nodes are hidden by a multithreaded overlapping scheme.

REFERENCES

[1] Xavier Álvarez-Farré, Andrey Gorobets, and F. Xavier Trias. A hierarchical parallel implementation for heterogeneous computing. Application to algebra-based CFD simulations on hybrid supercomputers. *Computers & Fluids*, 214:104768, 2021.