

A PORTABLE ALGEBRAIC IMPLEMENTATION FOR RELIABLE INDUSTRIAL LES

*M. Mosqueda-Otero¹, À. Alsalti-Baldellou^{1,2}, J. Plana-Riu¹, X.
Álvarez-Farré³, G. Colomer¹, A. Gorobets⁴, F. X. Trias¹, and A. Oliva¹*

¹ *Universitat Politècnica de Catalunya - BarcelonaTech, CTTC, Carrer de Colom 11, 08222 Terrassa, Spain*

² *University of Padova, Department of Information Engineering, Via Giovanni Gradenigo, 6b, 35131 Padova PD, Italy*

³ *High-Performance Computing Team - SURF, Science Park 140, 1098 XG Amsterdam, The Netherlands*

⁴ *Keldysh Institute of Applied Mathematics, 4A, Miusskaya Sq., Moscow 125047, Russia*

marcial.francisco.mosqueda@upc.edu

1 Introduction

The continuous development of novel numerical methods, coupled with the rapid evolution of high-performance computing (HPC) systems, has significantly expanded the role of computational fluid dynamics (CFD) in various industrial applications. Despite these advancements, the development of CFD faces persistent challenges. Early implementations were hindered by the compute-bound limitations of processors, which led to the adoption of compute-centric programming models. Over time, processor designs have evolved, addressing these limitations and resulting in a mismatch between computational power and memory bandwidth. This, in turn, forces the creation of complex memory hierarchies, complicating the optimization of traditional programs. At the same time, the widespread use of accelerators in diverse technological fields has driven the rise of hybrid architectures, offering greater computational throughput while improving power efficiency in large-scale applications (Witherden et al. 2014). However, this shift introduces a new challenge: ensuring the portability of legacy applications. This challenge requires versatile software architectures and the development of specialized APIs, such as CUDA, OpenCL, and HIP (Edwards et al. 2014; Zhang et al. 2013).

In this context, the conservative discretization method for unstructured grids, as proposed by Trias et al. (2014), has been adopted and implemented under TermoFluids Algebraic (TFA)—our in-house code based on an innovative algebra-dominant framework, HPC² (See Álvarez-Farré et al. 2018). This robust framework facilitates seamless integration into open-source codes and hybrid supercomputing environments, as shown by Komen et al. (2021).

While computational power has improved, the time and resources required for detailed simulations remain a significant bottleneck. Achieving large-scale simulations is crucial for meeting industry demands for rapid decision-making, shorter product development cycles, and expanding CFD’s industrial application

fields. Thus, our research seeks to ensure the integration of modern CFD methodologies into industry practices, enabling precise and accurate simulations of complex processes while efficiently utilizing available resources and reducing simulation costs within limited timeframes.

2 Portability for CFD

The construction of codes based on a minimal set of algebraic kernels has become essential for ensuring portability, optimization, and ease of maintenance, particularly in light of the growing diversity of computational architectures and hardware vendors. The hybrid nature of modern high-performance computing (HPC) systems presents additional challenges, as effective utilization of both processors and parallel accelerators often requires heterogeneous computations and complex data exchanges. Traditional CFD codes, however, rely on intricate data structures and specialized computational routines, which complicates portability. To address this, algorithms centered on algebraic kernels, such as the sparse matrix-vector product (SpMV), linear combination of vectors (axpy), element-wise product of vectors (axty), and the dot product, emerge as promising solutions (See Álvarez-Farré et al. 2018).

However, this approach introduces a computational challenge linked with the low arithmetic intensity of SpMV operations. This limitation can be mitigated using the more computationally intensive sparse matrix-matrix product (SpMM). Substituting SpMV with SpMM significantly reduces memory access demands and the memory footprint by reusing matrix coefficients (See Alsalti-Baldellou et al. 2023b).

Beyond portability, algebra-based CFD implementations offer distinct numerical benefits. For example, they facilitate the development of efficient CFL-like conditions, reducing computational cost by up to 4x compared to classical approaches (See Trias et al. 2023). Additionally, algebraic frameworks allow for the streamlined incorporation of advanced techniques





Figure 1: Strong scalability analysis; CPU-based system with $350 \times 480 \times 350$ - 58.8M CVs - grid comparing MPI-only vs. MPI+OpenMP paradigms (left plot) and GPU-based system with a $370 \times 400 \times 370$ - 54.76M CVs - grid (right plot)

such as flux limiters, yielding compact and efficient implementations by utilizing incidence matrices and local operations to control gradient ratios, thus reducing the number of computing kernels required for porting (See Valle et al. 2022).

Building on these principles, Alsalti-Baldellou et al. (2023a) proposed an effective approach to accelerate Poisson solvers by exploiting domain symmetries. By carefully ordering the unknowns, SpMV operations could be replaced with SpMM, leading to a 2.5x increase in the performance of compute-intensive kernels while significantly reducing the solver’s memory footprint and setup costs.

3 Algorithm scalability analysis

A numerical test evaluates the performance and scalability of TFA’s base algorithm under different parallel computing paradigms to ensure its capability to address large-scale problems. Specifically, we compare the performance of an MPI-only configuration—where each CPU core is assigned a single task—against a hybrid MPI+OpenMP configuration, which utilizes MPI processes and multi-threaded executions per node. Furthermore, the scalability of the code on hybrid HPC systems is analyzed, taking advantage of TFA’s underlying structure, which is based on minimal algebraic kernels. This architecture enables broad portability across diverse GPU hardware using OpenCL.

The numerical test case solves a turbulent channel flow using a conjugate gradient solver with a Jacobi preconditioner for Poisson’s equation, combined with an explicit time integration scheme and a variable time step. Since the primary objective is to assess the scalability of TFA+HPC² kernels, each case is limited to 10 time steps, with 800 solver iterations per step. The test problem is solved over the entire domain without exploiting symmetries, thereby isolating the key algebraic kernels —SpMV, axpy, axty, and dot product— for performance measurement and analysis.

CPU-based system tests were conducted using MPI-only and MPI+OpenMP configurations on the MareNostrum 5 GPP supercomputer at BSC. The experiments run on nodes equipped with two Intel Xeon Platinum 8480+ processors (56 cores, 2 GHz, 105 MB L3 cache, and 307.2 GB/s memory bandwidth) with 256 GB of RAM, interconnected via ConnectX-7 NDR200 InfiniBand. Additionally, GPU-accelerated tests were conducted on the MareNostrum 5 ACC supercomputer at BSC, using nodes with two Intel Xeon Platinum 8460Y processors (40 cores, 2.63 GHz, 105 MB L3 cache, and 307.2 GB/s memory bandwidth), 512 GB of RAM, four NVIDIA Hopper H100 with 64GB HBM2, and interconnected through four ConnectX-7 NDR200 InfiniBand cards.

Figure 1 illustrates the strong scalability results for both HPC architectures: (i) CPU-based system (left plot) applying a 1-node baseline with a $350 \times 480 \times 350$ grid and (ii) GPU-based system (right plot) using a 1-node baseline with a $370 \times 400 \times 370$ grid. The results show a marked super-linear speedup for hybrid MPI+OpenMP processes, which can be attributed to enhanced cache utilization; moreover, the MPI-only solution exhibits significant communication overhead, particularly when using 16 or more computational nodes. In addition, the strong scalability for TFA+HPC² on a hybrid architecture presents a steady speed-up as the problem scales. However, despite the memory-bound nature of CFD applications, the efficiency remains between 50% and 70% up to 16 nodes. For a 32-node implementation, the efficiency drops to approximately 40%, indicating the increasing impact of communication overhead in hybrid architectures.

On regards to the weak scalability analysis¹, Figure 2 presents both CPU-based (focus on MPI+OpenMP results) and GPU-based tests. The hybrid parallel

¹The main objective of this study is to assess the scalability of TFA+HPC² kernels. Thus, the number of iterations per time step was fixed to ensure consistent kernel calls as the problem scales.




Figure 2: Weak scalability analysis; MPI+OpenMP paradigm with 525k CVs per CPU-core, up to 128 nodes (left plot) and GPU-based implementation with 13.69M CVs per GPU card, up to 64 nodes (right plot)

paradigm (left plot) shows a 13% drop in performance when scaling up to 128 nodes and a 9% drop when scaling up to 64 nodes, referenced to the 1-node baseline. Similarly, the GPU-accelerated weak scalability (right plot) displays a 9% drop in performance when scaling up to 64 nodes, referenced to the 1-node baseline.

4 Closing remarks & Future work

TFA+HPC² framework demonstrates strong portability across various HPC architectures, leveraging its minimal algebraic kernel design to ensure broad compatibility and efficiency. Additionally, scalability analysis demonstrates that TFA+HPC² kernels maintain high scalability and computational efficiency across different HPC architectures, making them suitable for large-scale applications.

This work will be expanded by executing performance analysis, such as evaluating the roofline model. In addition, verification of our in-house CFD implementation (TFA+HPC²) will be included via the solution of different industrial cases.

Acknowledgments

This work is supported by the Ministerio de Economía y Competitividad, Spain, SIMEX project (PID2022-142174OB-I00). M.MO. is supported by the Catalan Agency for Management of University and Research Grants (2024 FI-1 00684). In addition, J.PR is also supported by the Catalan Agency for Management of University and Research Grants (2022 FI-B 00810). Calculations were performed on the MareNostrum 5 GPP and ACC supercomputers. The authors thankfully acknowledge these institutions.

References

Alsalti-Baldellou, À., X. Álvarez-Farré, F. X. Trias, and A. Oliva (2023a). “Exploiting spatial symmetries for solving Poisson’s equation”. In: *Journal of Computational Physics* 486, p. 112133.

Alsalti-Baldellou, À., G. Colomer, J. A. Hopman, X. Álvarez-Farré, A. Gorobets, F. X. Trias, C. D. Pérez-Segarra, and A. Oliva (2023b). “Reliable overnight industrial LES: challenges and limitations. Application to CSP technologies”. In: *14th International ERCOFTAC Symposium on Engineering, Turbulence, Modelling and Measurements: 6th-8th September 2023, Barcelona, Spain: proceedings*.

Álvarez-Farré, X., A. Gorobets, F. X. Trias, R. Borrell, and G. Oyarzun (2018). “HPC² – A fully portable algebra-dominant framework for heterogeneous computing. Application to CFD”. In: *Computers & Fluids* 173, pp. 285–292.

Edwards, H. C., C. R. Trott, and D. Sunderland (2014). “Kokkos: Enabling manycore performance portability through polymorphic memory access patterns”. In: *J. Parallel Distrib. Comput.* 74.12, pp. 3202–3216.

Komen, E., J. A. Hopman, E. M. A. Frederix, F. X. Trias, and R. W. C. P. Verstappen (2021). “A symmetry-preserving second-order time-accurate PISO-based method”. In: *Computers & Fluids* 225, p. 104979.

Trias, F. X., X. Álvarez-Farré, À. Alsalti-Baldellou, A. Gorobets, and A. Oliva (2023). “An Efficient Eigenvalue Bounding Method: Cfl Condition Revisited”.

Trias, F. X., O. Lehmkuhl, A. Oliva, C. D. Pérez-Segarra, and R. W. C. P. Verstappen (2014). “Symmetry-preserving discretization of Navier-Stokes equations on collocated unstructured meshes”. In: *J. Comput. Phys.* 258 (1), pp. 246–267.

Valle, N., X. Álvarez-Farré, A. Gorobets, J. Castro, A. Oliva, and F. X. Trias (2022). “On the implementation of flux limiters in algebraic frameworks”. In: *Computer Physics Communications* 271, p. 108230. ISSN: 0010-4655.

Witherden, F. D., A. M. Farrington, and P. E. Vincent (2014). “PyFR: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach”. In: *Comput. Phys. Commun.* 185.11, pp. 3028–3040.

Zhang, Y., M. Sinclair, and A. A. Chien (2013). “Improving Performance Portability in OpenCL Programs”. In: *Supercomputing*, pp. 136–150. ISBN: 978-3-642-38750-0.