ParCFD2022
33^{rd} International Conference on Parallel Computational Fluid Dynamics
May 25-27 2022, Alba Italy

DNS/LES USING A MINIMAL SET OF ALGEBRAIC KERNELS: CHALLENGES AND OPPORTUNITIES

F.X.Trias[1], X.Álvarez-Farré[1], A.Alsalti-Baldellou[1], A.Gorobets and A.Oliva[1]

[1] Heat and Mass Transfer Technological Center, Technical University of Catalonia, C/Colom 11, 08222 Terrassa (Barcelona), francesc.xavier.trias@upc.edu

Key words: DNS, LES, symmetry-preserving, portability, SpMM, eigenbounds

Abstract. In the last decades, CFD has become a standard design tool in many fields such as automotive, aeronautical and wind power industries. The driven force behind this is the development of numerical techniques in conjunction with the progress of high performance computing (HPC) systems. However, nowadays we can say that its legacy from the 90-2000s is hindering its progress. The reasons are two-fold: (i) codes designed for CPUs cannot be easily ported and optimized to new architectures (GPUs, ARM...) and (ii) legacy numerical methods chosen to solve (quasi)steady problems using RANS models are not appropriate for more accurate (and more expensive) techniques such as LES or DNS. This work aims to interlace these two pillars with the final goal to enable LES/DNS of industrial applications to be efficiently carried out on modern HPC systems while keeping codes easy to port and maintain. In this regard, a fully-conservative discretization for collocated unstructured grids is used: namely, it exactly preserves the symmetries of the underlying differential operators and is based on only five discrete operators (*i.e.* matrices): the cell-centered and staggered control volumes (diagonal matrices), Ω_c and Ω_s, the face normal vectors, N_s, the cell-to-face interpolation, $\Pi_{c\to s}$ and the cell-to-face divergence operator, M. Therefore, it constitutes a robust approach that can be easily implemented in already existing codes. Then, for the sake of cross-platform portability and optimization, CFD algorithms must rely on a very reduced set of (algebraic) kernels (*e.g.* sparse-matrix vector product, `SpMV`; dot product; linear combination of vectors). This imposes restrictions and challenges that need to be addressed such as the inherent low arithmetic intensity of the `SpMV`, the reformulation of flux limiters and boundary conditions or the efficient computation of eigenbounds to determine the time-step. Results showing the benefits of symmetry-preserving discretizations will be presented together with novel methods aiming to keep a good balance between portability and performance.

1 INTRODUCTION

The essence of turbulence are the smallest scales of motion [1]. They result from a subtle balance between convective transport and diffusive dissipation. Mathematically, these terms are governed by two differential operators differing in symmetry: the convective operator is skew-symmetric, whereas the diffusive is symmetric and negative-definite.

At discrete level, operator symmetries must be retained to preserve the analogous (invariant) properties of the continuous equations [2, 3]: namely, the convective operator is represented by a skew-symmetric coefficient matrix, the diffusive operator by a symmetric, negative-definite matrix and the divergence is minus the transpose of the gradient operator. Therefore, even for coarse grids, the energy of the resolved scales of motion is convected in a stable manner, *i.e.* the discrete convective operator transports energy from a resolved scale of motion to other resolved scales without dissipating any energy, as it should be from a physical point-of-view. Furthermore, high-order symmetry-preserving discretizatons can be constructed for Cartesian staggered grids [2]. It is noteworthy to mention that in the last decade, many DNS reference results have been successfully generated using this type of discretization (see Figure 1, for example).

However, for unstructured meshes, it is (still) a common argument that accuracy should take precedence over the properties of the operators. Contrary to this, our philosophy is that operator symmetries are critical to the dynamics of turbulence and must be preserved. With this in mind, a fully-conservative discretization method for general unstructured grids was proposed in Ref. [3]: it exactly preserves the symmetries of the underlying differential operators on a collocated mesh. In summary, and following the same notation than in Ref. [3], the method is based on a set of five basic operators: the cell-centered and staggered control volumes (diagonal matrices), Ω_c and Ω_s, the matrix containing the face normal vectors, N_s, the cell-to-face scalar field interpolation, $\Pi_{c\to s}$ and the cell-to-face divergence operator, M. Once these operators are constructed, the rest follows straightforwardly from them. Therefore, the proposed method constitutes a robust and easy-to-implement approach to solve incompressible turbulent flows in complex configurations that can be easily implemented in already existing codes such as OpenFOAM® [4].

2 RETHINKING CFD FOR PRESENT AND FUTURE PORTABILITY

Building codes on top of a minimal set of basic kernels is the cornerstone for portability and optimization. This became even more crucial due to the increasing variety of computational architectures competing in the exacale race. Moreover, the hybridization of HPC systems imposes additional constraints, since heterogeneous computations are needed to efficiently engage processors and massively-parallel accelerators. This involves different parallel paradigms and computing frameworks and requires complex data exchanges between computing units. However, (legacy) CFD codes usually rely on sophisticated data structures and computing subroutines, making portability terribly cumbersome.

In this context, we proposed [6] a completely different approach: for the sake of cross-platform portability and optimization, CFD algorithms must rely on a very reduced set of (algebraic) kernels (*e.g.* sparse-matrix vector product, `SpMV`; dot product; linear combination of vectors). This imposes restrictions and challenges that need to be addressed such as the inherent low arithmetic intensity (AI) of the `SpMV`, the reformulation of flux limiters [7] and boundary conditions or the efficient computation of eigenbounds to determine the time-step, Δt.

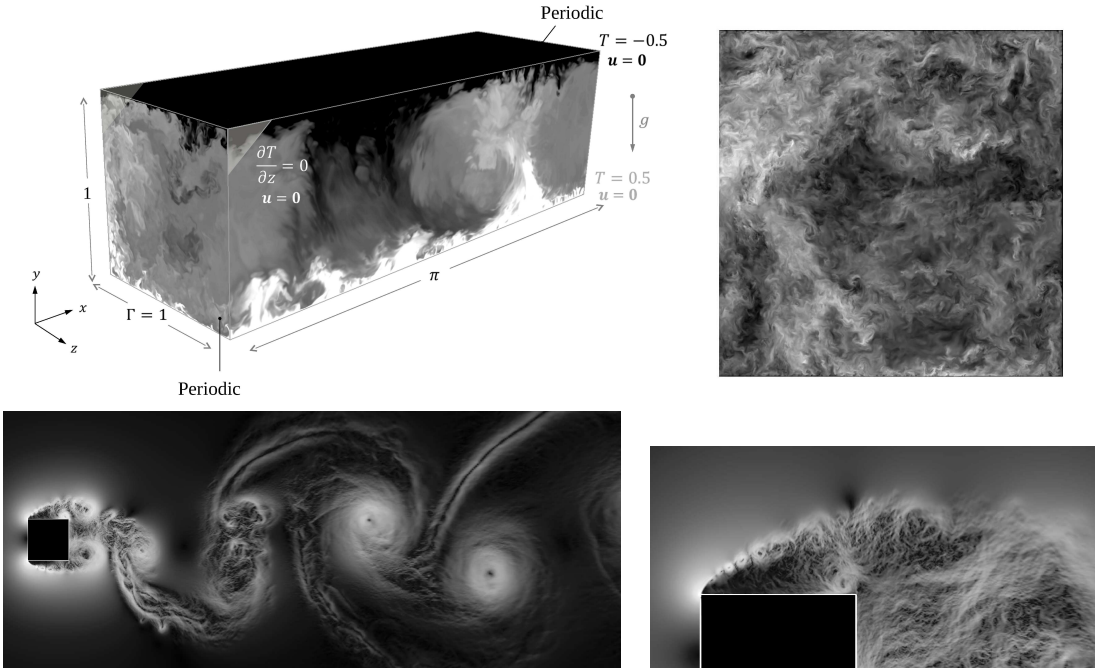

Figure 1: Examples of DNSs computed using symmetry-preserving discretizations. Top: air-filled ($Pr = 0.7$) Rayleigh-Bénard configuration studied in Ref. [1]. Instantaneous temperature field at $Ra = 10^{10}$ (left) and instantaneous velocity magnitude at $Ra = 10^{11}$ (right) for a spanwise cross section are shown. The latter was computed on 8192 CPU cores of the MareNostrum 4 supercomputer on a mesh of 5.7 billion grid points. Bottom: DNS of the turbulent flow around a square cylinder at $Re = 22000$ computed on 784 CPU cores of the MareNostrum 3 supercomputer on a mesh of 323 million grid points [5] .

3 CHALLENGES AND OPPORTUNITIES

Relying on a very reduced set of algebraic kernels enables code portability and facilitates its maintenance and optimization. However, it comes together with two types of challenges and restrictions. Firstly, *computational challenges* such as the inherent low AI of the `SpMV`. This can be strongly aliviated by using the sparse matrix-matrix product, `SpMM`, which has a much higher AI and, therefore, performance. This can be done in a great variety of situations: *e.g.* multiple transport equations, cases with spatial reflection symmetries, parallel-in-time simulations and, in general, in any situation where a matrix, $\hat{\mathsf{A}} \in \mathbb{R}^{N\times N}$, can be viewed as a Kronecker product between a diagonal matrix $\mathsf{C} \equiv diag(\boldsymbol{c}) \in \mathbb{R}^{K\times K}$ and a sparse matrix, $\mathsf{A} \in \mathbb{R}^{N/K\times N/K}$, *i.e.*

$$\boldsymbol{y} = \hat{\mathsf{A}}\boldsymbol{x} \text{ (with } \texttt{SpMV}) \qquad \{\boldsymbol{y}_1, \ldots, \boldsymbol{y}_K\} = \mathsf{A}\,\{c_1\boldsymbol{x}_1, \ldots, c_K\boldsymbol{x}_K\} \text{ (with } \texttt{SpMM}), \tag{1}$$

where $\hat{\mathsf{A}} = \mathsf{C}\otimes\mathsf{A}$ and $\boldsymbol{x}_i, \boldsymbol{y}_i \in \mathbb{R}^{N/K}$. In this way, matrix coefficients are re-used leading to a significantly higher AI and lower memory consumption. Details about the implementation and performance analysis of the `SpMM` on different architectures and its impact on the solution of Poisson's equation will be presented in two companion works [8, 9].

Secondly, *algorithmic challenges* such as the reformulation of classical flux limiters [7] or the boundary conditions also have to be addressed. The latter can be naturally solved

by casting boundary conditions into an affine transformation

$$\boldsymbol{\varphi}_h \rightarrow \mathsf{A}\boldsymbol{\varphi}_h + \boldsymbol{b}_h, \tag{2}$$

which allows a purely algebraic treatment of virtually all existing boundary conditions for both explicit and implicit time-integration methods. Furthermore, a new accurate and portable (only based on the above explained algebraic kernels) approach *à la CFL* for bounding the eigenvalues of the convective and diffusive operators will be presented. Results will be presented during the conference.

Acknowledgments. F.X.T., X.A-F, A.A-B and A.O. are supported by the *Ministerio de Economía y Competitividad*, Spain, RETOtwin project (PDC2021-120970-I00). F.X.T. and A.O. are also supported by the *Generalitat de Catalunya* RIS3CAT-FEDER, FusionCAT project (001-P-001722). A.A-B is also supported by the predoctoral grants DIN2018-010061 and 2019-DI-90, by MCIN/AEI/10.13039/501100011033 and the Catalan Agency for Management of University and Research Grants (AGAUR). Calculations were performed on the MareNostrum 4 supercomputer at the BSC. The authors thankfully acknowledge these institutions.

References

[1] F. Dabbagh, F. X. Trias, A. Gorobets, and A. Oliva. Flow topology dynamics in a three-dimensional phase space for turbulent Rayleigh-Bénard convection. *Physical Review Fluids*, 5:024603, 2020.

[2] R. W. C. P. Verstappen and A. E. P. Veldman. Symmetry-Preserving Discretization of Turbulent Flow. *Journal of Computational Physics*, 187:343–368, 2003.

[3] F. X. Trias, O. Lehmkuhl, A. Oliva, C.D. Pérez-Segarra, and R.W.C.P. Verstappen. Symmetry-preserving discretization of Navier-Stokes equations on collocated unstructured meshes. *Journal of Computational Physics*, 258:246–267, 2014.

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

[5] F. X. Trias, A. Gorobets, and A. Oliva. Turbulent flow around a square cylinder at Reynolds number 22000: a DNS study. *Computers & Fluids*, 123:87–98, 2015.

[6] X. Álvarez, A. Gorobets, F. X. Trias, R. Borrell, and G. Oyarzun. HPC^2 - a fully portable algebra-dominant framework for heterogeneous computing. Application to CFD. *Computers & Fluids*, 173:285–292, 2018.

[7] N. Valle, X. Álvarez-Farré, A. Gorobets, J. Castro, A. Oliva, and F. X. Trias. On the implementation of flux limiters in algebraic frameworks. *Computer Physics Communications*, 271:108230, 2022.

[8] X. Álvarez-Farré, À. Alsalti-Baldellou, A. Gorobets, A. Oliva, and F. X. Trias. On the benefits and applications of sparse matrix-matrix product on various parallel architectures. In *Parallel Computational Fluids Dynamics*, Alba, Italy, May 2022.

[9] À. Alsalti-Baldellou, X. Álvarez-Farré, A. Gorobets, A. Oliva, and F. X. Trias. Strategies to increase the arithmetic intensity of linear solvers. In *Parallel Computational Fluids Dynamics*, Alba, Italy, May 2022.