14th International ERCOFTAC symposium on engineering, turbulence, modelling and measurements (ETMM)
Barcelona, Spain, 6-8 September, 2023

Reliable overnight industrial LES: challenges and limitations. Application to CSP technologies

A. Alsalti-Baldellou[1,2], G. Colomer[1], J. A. Hopman[1], X. Álvarez-Farré[3], A. Gorobets[4], F. X. Trias[1] and A. Oliva[1]

[1] Heat and Mass Transfer Technological Center, Universitat Politècnica de Catalunya – BarcelonaTech, Carrer de Colom 11, 08222 Terrassa (Barcelona), Spain

[2] Termo Fluids SL, Sabadell (Barcelona), Spain; www.termofluids.com

[3] High-Performance Computing Team, SURF, Science Park 140, 1098 XG Amsterdam, Netherlands

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

Key words: DNS, LES, symmetry-preserving, unstructured grids, collocated

Abstract. Preserving the operators' symmetries at the discrete level is crucial to enable reliable DNS and LES simulations of turbulent flows. Moreover, real-world applications demand robust and stable numerical methods suitable for complex geometries. In this regard, this work presents a symmetry-preserving discretisation for unstructured collocated grids that, apart from being virtually free of artificial dissipation, is shown to be unconditionally stable. To ensure cross-platform portability, the implementation of such a discretisation relies on a minimal set of algebraic kernels. Doing this poses challenges that need to be addressed, like the low arithmetic intensity of the sparse matrix-vector product, the reformulation of boundary conditions and flux limiters, or the efficient computation of eigenbounds to determine the time-step. Finally, a comparison with open-source CFD codes will be made, and a relevant case from the CSP industry will be presented in order to assess the feasibility of overnight industrial simulations.

1 Introduction

In the last decades, CFD has become a standard design tool in many fields, such as the automotive, aeronautical, and renewable energy industries. The driving force behind this is the development of numerical techniques in conjunction with the progress of high-performance computing (HPC) systems. However, progress is nowadays hindered by its legacy from the 90-2000s. The reasons are two-fold. Firstly, the design of digital processors constantly evolves to overcome limitations and bottlenecks. The formerly compute-bound nature of processors led to compute-centric programming languages and simulation codes. However, raw computing power grows faster than the memory bandwidth, turning around the problem and leading to increasingly complex memory hierarchies that make optimising traditional applications a cumbersome task. Moreover, new parallel programming languages emerged to target modern hardware, *e.g.*, CUDA, HIP and oneAPI, and porting algorithms and applications has become restrictive. Secondly, legacy numerical methods chosen to solve (quasi)steady problems using RANS models are inappropriate for more accurate (and expensive) techniques such as LES or DNS. We aim to interlace these two pillars with the final goal of enabling LES and DNS of industrial applications to be efficiently carried out on modern HPC systems while keeping codes easy to port, optimise, and maintain. In this regard, the fully-conservative discretisation for collocated unstructured grids proposed in [1] is adopted: it constitutes a very robust approach that can be easily implemented in existing codes such as

OpenFOAM [2].

The main recognised limitations of LES in the industry are their computational cost and the wall-clock simulation time. Thanks to the above-explained advent of new computational architectures, the former is becoming less and less critical, whereas the latter is still the most limiting factor precluding LES from being routinely used in the industry. For that to be possible, the consensus is that widespread adoption in the industry begins when a run can be carried out overnight [3]. Namely, the industry is governed by shortening design cycles, faster time-to-market, and increased expectations of operability and reliability for established product lines. Therefore, it is willing to spend on hardware and software as long as analysts can obtain meaningful insights in a time commensurate with design cycles. Overnight runs fit this timescale, and in this context, we aim to answer the following question: *can we use LES modelling to simulate complex industrial flows with overnight simulations accurately?*

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 optimisation, which became crucial after the increasing variety of computational architectures competing in the exascale race. Moreover, the hybridisation of HPC systems imposes additional constraints since heterogeneous computations are usually needed to engage processors and massively parallel accelerators efficiently. 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 extremely complex.

In this context, we proposed a completely different approach. That is, making CFD algorithms rely on a very reduced set of algebraic kernels, *e.g.*, the sparse matrix-vector product (`SpMV`), the linear combination and the dot product of vectors [4]. This imposes restrictions and challenges that need to be addressed, such as the low arithmetic intensity of the `SpMV`, the reformulation of boundary conditions and flux limiters [5], or the efficient computation of eigenbounds to determine the time-step, Δt.

3 Challenges and opportunities

Relying on a minimal set of algebraic kernels enables code portability and facilitates its maintenance and optimisation. However, it comes together with two types of challenges and restrictions. Firstly, *computational challenges* like the low arithmetic intensity of the `SpMV`, which can be alleviated by using the more compute-intensive sparse matrix-matrix product (`SpMM`). This is possible in a great variety of situations, such as with multiple transport equations, in cases with spatial reflection symmetries, parallel-in-time simulations and, in general, whenever dealing with matrices, $\hat{\mathsf{A}} \in \mathbb{R}^{N\times N}$, decomposable as the Kronecker product of a diagonal matrix, $\mathsf{C} \equiv diag(c) \in \mathbb{R}^{K\times K}$, and a sparse matrix, $\mathsf{A} \in \mathbb{R}^{N/K\times N/K}$, *i.e.*, $\hat{\mathsf{A}} = \mathsf{C} \otimes \mathsf{A}$. Indeed, under such circumstances, the standard `SpMV` can be replaced with the `SpMM`:

$$y = \hat{\mathsf{A}}x \quad \Longrightarrow \quad (y_1, \ldots, y_K) = \mathsf{A}\,(c_1 x_1, \ldots, c_K x_K)\,, \tag{1}$$

where $x_i, y_i \in \mathbb{R}^{N/K}$. By doing so, matrix coefficients are recycled, thus significantly reducing the memory accesses and the memory footprint of the operators.

Secondly, *algorithmic challenges* such as reformulating classical flux limiters [5] or the boundary conditions must also be addressed. The latter can be naturally solved by casting boundary conditions into an affine transformation:

$$\varphi_h \to \mathsf{A}\varphi_h + b_h, \tag{2}$$

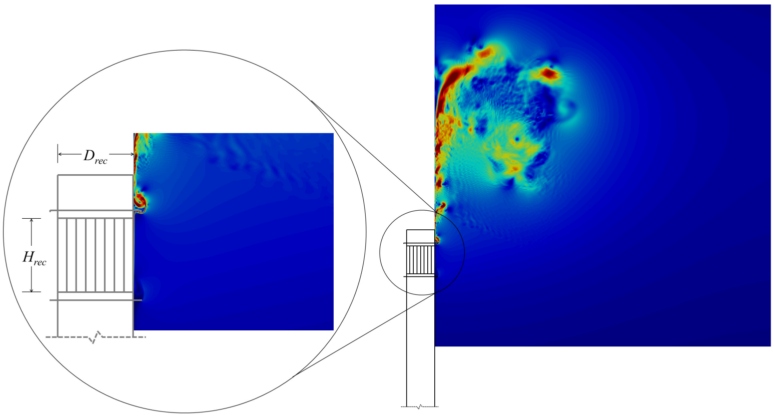

Figure 1: CSP simulation of a central tower receiver.

allowing a purely algebraic treatment of virtually all existing boundary conditions both for explicit and implicit time-integration methods. Furthermore, an accurate and portable approach solely relying on the above-mentioned algebraic kernels for bounding the eigenvalues of the convective and diffusive operators has also been proposed.

4 Performance analysis and application to CSP technologies

Performance analysis of the code will be done in two stages. Firstly, we will focus on the performance analysis for a given mesh without considering the results' accuracy The comparison will be made with open-source CFD codes like OpenFOAM and Code Saturne. Several factors will be analysed separately and compared, such as exploiting the shared-memory paradigm with OpenMP, increasing the arithmetic intensity through the strategies mentioned above, and the effect of GPUs. In the second stage, we will analyse those aspects, apart from the mesh resolution, that can affect the accuracy of the results. Comparison will be made with best practices on the above-mentioned open-source CFD codes. Several factors will be analysed separately and compared. Namely, the spatial and temporal discretisation, the Poisson solver and the turbulence models. Finally, a relevant case from the CSP industry (see Figure 1) will be studied as a demonstrative test case to assess the feasibility of overnight industrial simulations.

REFERENCES

[1] 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.

[2] 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.

[3] Rainald Löhner, Carsten Othmer, M. Mrosek, Alejandro Figueroa, and Atis Degro. Overnight industrial les for external aerodynamics. *Computers & Fluids*, 214:104771, 2021.

[4] X. Álvarez Farré, 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.

[5] 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.