HORIZON-EUROHPC-JU-2021-COE-01

EXCELLERAT P2

The European Centre of Excellence for Engineering Applications

Project Number: 101092621

D2.10 Report on the Alya Application Use Case

The EXCELLERAT P2 project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101092621. The JU receives support from the European Union's Horizon Europe research and innovation programme and Germany, Italy, Slovenia, Spain, Sweden and France.

Work Package:	2	Use-Case Execution
Author(s):	Herbert Owen	BSC
Approved by	Executive Centre Management	17.11.2023
Reviewer	Guillaume Daviller	Cerfacs
Reviewer	Fulvio Stella	URMLS
Dissemination Level	Public	

Date	Author	Comments	Version	Status
26.10.2023	Herbert Owen	First draft	V0.1	Draft
2.11.2023	Herbert Owen	Second draft	V0.2	Draft
14.11.2023	Herbert Owen	Third draft	V0.3	Draft
16.11.2023	Herbert Owen	Final version	V1.0	Final

List of abbreviations

AI	Artificial Intelligence
AFC	Active Flow Control
CAD	Computer Aided Design
CFD	Computational Fluid Dynamics
CO2	Carbon Dioxide
CRM	Common Research Model
DMD	Dynamic Mode Decomposition
DNS	Direct Numerical Simulation
DoF	Degree of Freedom
DRL	Deep Reinforcement Learning
GPU	Graphics Processing Unit
HDF5	Hierarchical Data Format version 5
HLPW	High Lift Prediction Workshop
HPC	High Performance Computing
I/O	Input Output
NOx	Nitrous Oxide
POD	Proper Orthogonal Decomposition
UC	Use Case
VR	Virtual Reality
WMLES	Wall Modelled Large Eddy Simulation
WP	Work Package

Executive Summary

This document summarizes the progress made in the 4th use case (UC-4), Fully integrated aircraft simulations with emissions models during the first year of EXCELLERAT P2.

The first section introduces the use case, and the second details its objectives. The workflow is presented in the third section. The fourth section describes progress made during the first year, and section five presents the steps for the following period.

In summary, the workflow development for UC-4 has progressed according to the schedule defined in deliverable D2.1. Work has been performed on the individual tasks planned for the project's first year.

This use case focuses on generating a multiphysics platform that integrates advanced simulation techniques from Computational Fluid Dynamics (CFD) codes and massively parallel workflows using leading-edge HPC architectures. Fully integrated simulations of the aircraft and propulsive system will be conducted with advanced High Performance Computing (HPC) algorithms to obtain detailed predictions of aerodynamic parameters and emissions characteristics, including non-CO2 effects like nitrous oxide (NOx) or soot.

To enable digital twins of fully integrated aircraft simulations, we plan to integrate coupled external aerodynamics with engine emissions simulations over scale resolving simulation methods, including flow control devices and high-fidelity combustion simulations.

In EXCELLERAT P1, Alya's standard low-order finite element version was used. For EXCELLERAT P2, we are working on a new high-order version. During the project's first year, we focused on validating the code for problems of increased complexity. Section four presents result for a Taylor Green vortex, channel flow, flow around a Windsor model car, and a simplified aircraft. Progress has been slightly slower than expected due to a lack of computational resources originated by the delay in the installation of Marenostrum 5. We were expecting it would be ready by July 2023, but it has been postponed until February 2024.

The project's second year will be devoted to improving Wall modelled Large Eddy Simulations (WMLES) of more complex geometries and coupling the flow simulations with a combustion model. For the second task, the Combustion group, led by Daniel Mira, has an open position we hope to fill soon.

Table of Contents

Table of Figures

1 Introduction

The fully integrated aircraft simulations with emissions models use case is aligned with the strategies in the aeronautical industry to achieve the EU decarbonisation objectives by 2050 in air transportation, which includes substantial reductions in CO2, NOx emissions, and noise as defined in the Flightpath 2050. Generating digital tools that can be used to obtain detailed information on the relevant parameters from the aircraft and propulsive systems is an essential step toward designing low-emissions ultra-silent aircraft. An efficient strategy to develop these digital tools is using data-driven methods based on deep learning, where a neural network is fed with data from experiments and/or high-fidelity simulations. These tools can be integrated into design and optimisation workflows to reduce the computational cost, turnaround time, and energy consumption during the design phases.
This use case focuses on generating a multiphysics platform that integrates advanced simulation techniques from CFD codes and massively parallel workflows using leading-edge HPC architectures. Fully integrated simulations of the aircraft and propulsive system will be conducted with advanced HPC algorithms to obtain detailed predictions of aerodynamic parameters and emissions characteristics, including non-CO2 effects like NOx or soot.
Figure 1 shows the flow field around the Common Research Model (CRM), a simplified aircraft geometry developed to test the accuracy of numerical simulations. Alya has already participated in previous versions of the High Lift Prediction Workshop (HLPW). It is currently active in the latest version (https://hiliftpw.larc.nasa.gov), where we expect to showcase the improvements developed within EXCELLERAT P2.

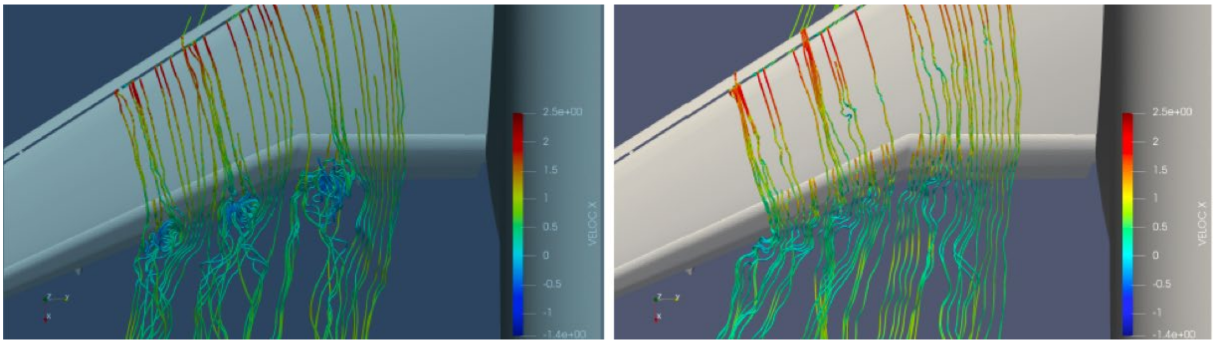

Figure 1: Instantaneous vortical structures calculated with Alya.

2 Objectives of the Use Case

The accomplishment of fully integrated simulations of the aircraft and propulsive system involves different levels of hierarchies and parallelism in the algorithms responsible for controlling the different tasks. This requires the revision of the load balancing strategies for inter- and intra-nodes in the CFD codes, but also communication with the other tasks during post-processing and data analytics. A multicode strategy will be used to perform coupled multiphysics simulations, which brings additional challenges to ensure high computational performance, flexibility, and portability. As Alya contains a fully parallel workflow, including the reading and partitioning, a fully parallel workflow can be achieved after the mesh generation step. Such a parallel workflow is required in a pre-processing step. The I/O will be adapted to facilitate the reading/writing of files during runtime, and parallel algorithms will be used to communicate the data.
To enable digital twins of fully integrated aircraft simulations, during the four years of the project we plan to integrate coupled external aerodynamics with engine emissions simulations over scale resolving simulation methods, including:

1. flow control devices (Active Flow Control (AFC), active surfaces, etc.).

2. high-fidelity combustion simulations with emissions.
3. Deep reinforced learning for flow control of high lift wings.

In EXCELLERAT P1, Alya's standard low-order finite element version was used. However, for EXCELLERAT P2, we are working on a new high-order version. It is based on a spectral element method of arbitrary order that uses Lobatto-Gauss-Legendre quadrature and close rule integration. For the convective term a skew-symmetric operator split is used to counter undesired aliasing effects generated by closed rule integration [1]. An entropy viscosity stabilisation is used for shock capturing [2]. Implicit and explicit time discretisation's are available. The use of OpenACC provides a code that is valid for both CPU & GPU.
During the first year of the project the main objective has been to develop and validate the high order version of Alya that takes full advantage of GPUs on compressible flow problems. For the second year our main objective will be related to the coupling with high-fidelity combustion simulations, but we will try to start working on flow control too.

3 Workflow Description

The workflow to be executed with Alya and its auxiliary tools is schematically depicted in Figure 2.

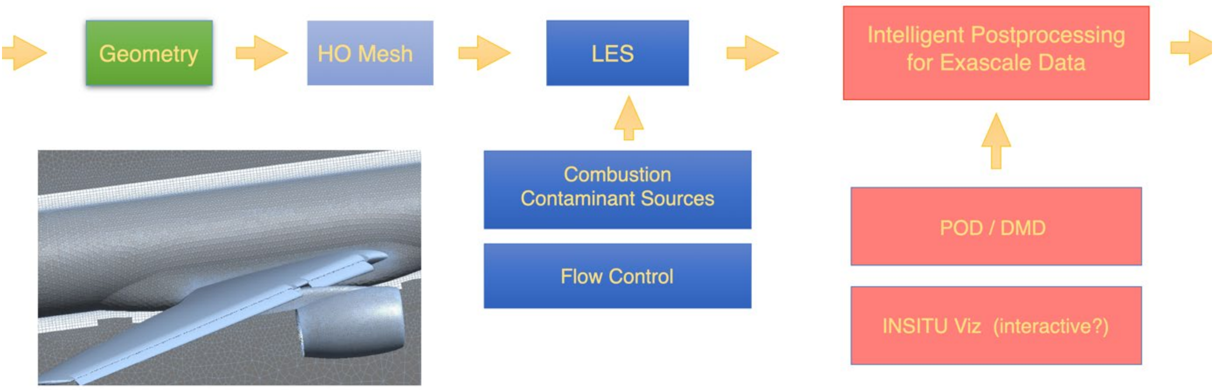

Figure 2: Workflow for fully integrated aircraft simulations with emissions models

The workflow starts with a complex geometry provided in some CAD format. In the case of the High Lift Prediction Workshop (HLPW) [6], the geometry is provided by the Workshop organiser. Before performing the simulation, a mesh must be obtained. We typically rely on commercial tools for complex geometries, ANSA [8] being our default option. However, we have also been experimenting with the open-source tool GMSH [4] for high-order meshes. For HLPW [6], the organisers also provide suggested grids to promote a better comparison between the different codes.

Wall-modelled Large Eddy Simulation (WMLES) lies at the heart of our workflow. Accuracy is still a challenge for high Reynolds number turbulent flows over complex geometries. The substantial computational cost means high-fidelity numerical methods must be used optimally in the most energy-efficient architectures. We are opting for a spectral high-order discretisation implemented optimally on GPUs. As the workflow indicates, the flow simulation is coupled to contaminant sources from a combustion model also available in Alya. The wall model we currently use for high-order elements closely follows what we have used for some time with low-order elements [11]. The first ingredient is the usage of average velocities as an input for the wall model. The second ingredient is that we can obtain the input velocity at an arbitrary distance from the wall that does not need to coincide with the first off wall point.

The Alya team has experience in active methods to control turbulence and separation, a crucial aspect of achieving high aerodynamic efficiency. Moreover, data-driven methods and their application to flow control will be explored with a focus on Deep Reinforcement Learning (DRL).

The ever-growing computational power of modern HPC systems enables us to perform detailed simulations of the flow, which give us access to massive information on the flow, however, typically only a tiny portion of the information is used. In an extreme case, even only the drag and lift coefficients are evaluated. The amount of data that is extracted depends heavily on the experience of the CFD user. Thus, for exascale simulations, intelligent postprocessing tools help to analyse and understand the flow making the workflow more efficient. Dynamic Mode Decomposition (DMD) and Proper Orthogonal Decomposition (POD) tools can help reduce the complexity of the data and identify the underlying physical processes. DMD decomposes the data into a set of modes that capture the dominant dynamics of the system. It helps to identify the frequencies and amplitudes of the dominant oscillations in the system and provides insight into its stability and behaviour. POD decomposes the data into a set of orthogonal modes that capture the dominant spatial structures in the data. It can be used to identify the most important patterns of variability in the system and help identify regions of high turbulence and flow separation. Using DMD and POD together can provide a complete understanding of the system by identifying the dominant modes of variability and the dominant spatial structures. BSC develops its own tools to use DMD/POD in HPC systems.

Additionally, we will use in situ visualisation and Virtual Reality (VR) to further enhance the analysis of exascale CFD data. Alya will be interfaced with the Distributed Data-parallel Scientific Visualisation tool VISTLE [7]. In situ visualisation helps analyse and visualise the data while it is being generated, rather than waiting until the simulation is complete, making our workflow more efficient. Additionally, it reduces the amount of data that needs to be stored and transferred, which is critical at Exascale. On the other hand, VR capabilities provided by VISTLE allow a more immersive and intuitive way to visualise and explore the data. This can be particularly useful for visualising complex, multi-dimensional datasets like the ones we generate. A combination of DMD, POD, in situ visualisation, and VR will help us transform how we analyse and visualise the data we generate.

We resume the key feature of each of the components of our workflow:

Mesh Generation

- Unstructured high-order tetrahedral grids generated with commercial (ANSA [8]) or free (GMSH [4]) tools.
- Partitioning to hexahedral third order meshes in GMSH and improvement with HOPR [5].
- In-house mesh partitioner that uses the parallel mesh partitioning library GeMPa [10] based on Space-Filling Curves.
- Using high memory nodes, we expect to generate O(10^9) nodes.

LES/CFD Simulation

- Fully parallel workflow starting from an already partitioned grid using external tools.
- O(10^9) degrees of freedom (DoF). Higher number of DoFs either using other external mesh generation tools to be investigated, automatic mesh subdivision, or higher order elements.
- Perform simulation until the time averaged values of interest are sufficiently converged.

- Medium sized job*: O(100) GPUs, O(96 hours).
- Instantaneous and averaged quantities passed to intelligent post processing tools using parallel output (Hierarchical Data Format version 5 (HDF5)) or in situ.
- For contaminants dispersion coupling with combustion module as a source.

Reactive flow

- Extension of the compressible Navier-Stokes solver based on spectral elements to turbulent reactive flows.
- Study of thermoacoustic instabilities of the PRECCINSTA burner from the high-pressure combustion test rig HIPOT at DLR Stuttgart.
- Expected runs with highly resolved LES with O(10^9) degrees of freedom.
- Job size: O(100) GPUs, O(48 hours).
- Combine with parallel post-processing/analysis using POD/DMD.

Flow control

- Passive methods based on modifying the airflow by means of geometrical additions on the shape.
- Active methods based on directly modifying the flow (e.g., by blowing-sucking means, etc.).
- Used for drag reduction/lift enhancement (aerodynamic efficiency) in aerospace applications.
- Classic methods (based on linear control theory) vs new generation methods based on AI (deep reinforcement learning, etc.).
- Several possible actuators are available, including active flow control, active surfaces, and deep reinforced learning for flow control of high lift wings.
- How to sense the flow for the actuation is still an open issue.
- Job size: O(10) GPUs, O(48 hours).

Intelligent postprocessing

- Parallel POD/DMD tools.
- In situ visualisation with Vistle [7] with the possibility of Virtual Reality.

To summarise, an efficient workflow for scale-resolving simulation of aircraft simulations with emissions models is proposed. Apart from the mesh generation, the workflow is fully parallel. A high-order discretisation is used for the LES simulation. It is implemented to work optimally on GPUs leading to a significant reduction in energy consumption.

According to their sizes, the runs that have been planned for EXCELLERAT P2 can be classified as follows:

- Hero run: coupled external aerodynamics with engine emissions simulations.
- Large-scale: External aerodynamics with scale-resolving simulation methods, including flow control devices (AFC, active surfaces, etc.), high-fidelity combustion simulations to predict pollutant formation NOx and soot.
- Small scale/ensemble: Deep Reinforce Learning for flow control of high lift wings.

4 Progress achieved

As mentioned in section 2, we are working on a new high-order version of Alya that provides superior accuracy than the low-order version used in EXCELLERAT P1. Therefore, during the project's first year, we focused on validating the code on problems of increased complexity.
To test that the code introduces minimal numerical dissipation, which is key for accurate large eddy simulation of turbulent flows, the Taylor Green Vortex has been analysed at two different Mach numbers, Ma=0.1 and Ma=1.25. It is a well-known problem with an exact closed-form solution of the incompressible Navier–Stokes equations involving the unsteady flow of a decaying vortex. The domain is a cube with periodic boundary conditions. As the vortex decays, smaller structures are formed, requiring numerical methods with minimal dissipation for a correct resolution. Figure 3 shows the evolution of the vortex for different time steps.

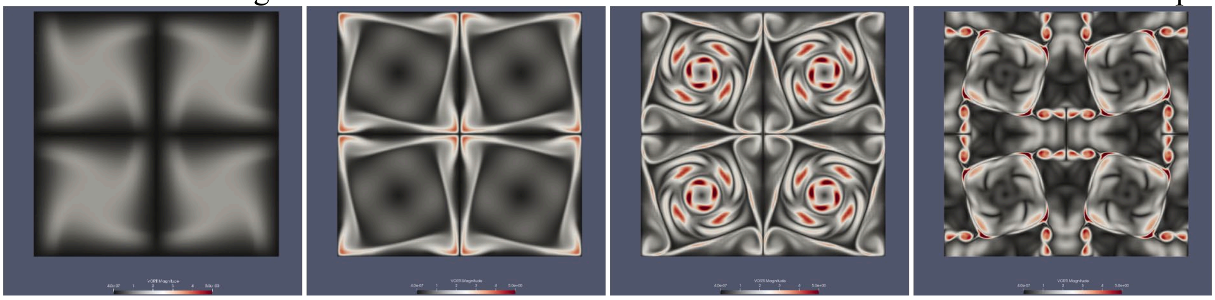

Figure 3: Flow vorticity at times 2, 7, 15 and 20s.

Figure 4 compares the temporal evolution of the kinetic energy and total dissipation obtained with Alya against a reference Direct Numerical Simulation (DNS) result. With the coarse mesh, some dissipation is observed but results are already much better than those that been obtained in the past with the low order version of Alya used in EXCELLERAT P1. Both medium and fine meshes provide excellent results with a clear convergence towards the DNS solution.

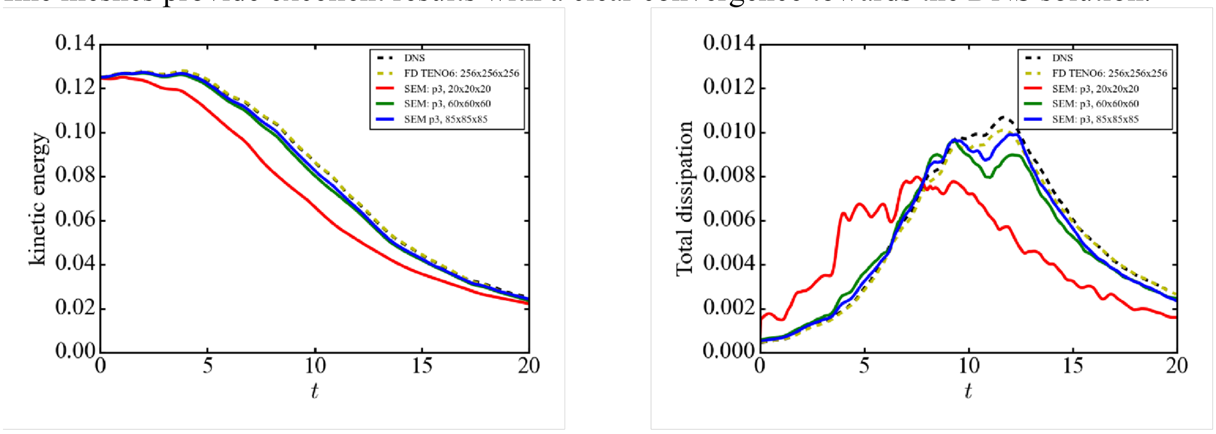

Figure 4: Temporal evolution of the kinetic energy and total dissipation with three different meshes using Alya compared to a Finite Difference solution and a reference DNS result.

The code has also provided excellent results for a channel flow at Reynolds number, Re_tau=950 and the turbulent flow behind a cylinder, which are omitted here due to space limitations.
Since the objective of Alya is the simulation of problems of engineering interest, we have lately concentrated on more complex problems, such as the Windsor model car, which is one of the two Benchmark cases for the Fourth Automotive Workshop – AutoCFD 2024 [9]. A mesh with 150 million DoFs, shown in Figure 5, has been used. The flow structures formed around the car are shown in Figure 6.

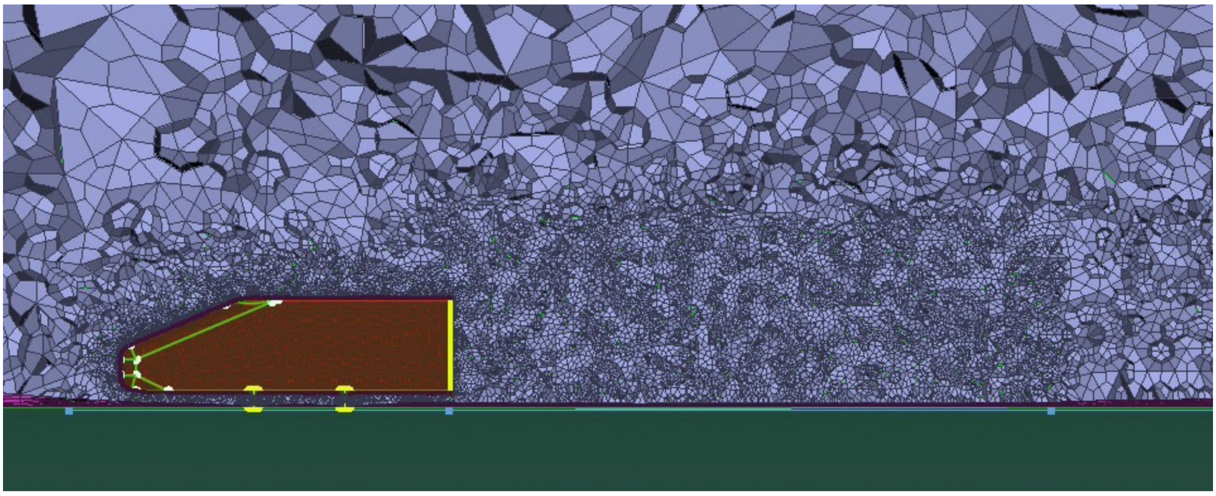

Figure 5: Windsor model car mesh.

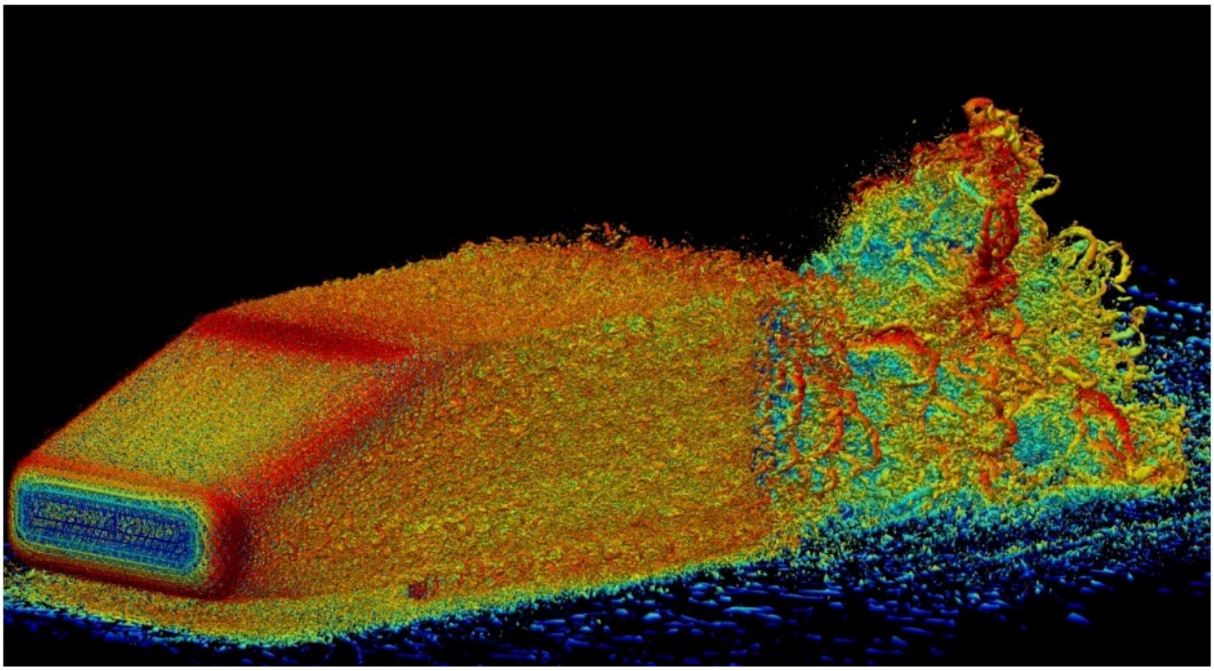

Figure 6: Q-criterion showing flow structures.

Finally, we are working on the high-order hexahedral mesh for the simplified CRM aircraft for the Fifth High Lift Prediction workshop. We have simulated Case 1 of the HLPW using slip boundary conditions. Until recently, we were experiencing robustness issues when WMLES was used on an aircraft, but it worked fine on simpler geometries. During the last weeks, we have improved the discretisation of the convective term following [3]. The latest modifications have significantly improved the robustness of WMLES for complex cases. We are now being able to run the CRM case satisfactorily. Figure 7 shows the velocity magnitude over the CRM wing that was obtained a few days before finishing this report. Advance has been slow because both Marenostrum 4 and MN-CTE-Power are saturated, especially for big cases. We will advance faster once Marenostrum 5 is available. While our code can take full advantage of GPUs, we are being forced to run on the CPUs due to the delays in installing Marenostrum 5. We had planned our work, supposing that Marenostrum 5 would be ready by July 2023, but now it seems we will have it by February 2024. We have asked for resources in LEONARDO, indicating our runs would typically require between 100 and 200 Nvidia A100 GPUs. Our proposal was rejected because they expected much bigger runs of the order of 2000 GPUs. We hope to run such simulations by the end of the project, but during the first years, we need to start more progressively. High order meshing of complex geometries has also posed some challenges. We initially tested the open-source tools GMSH [4], HOPR [5], and the commercial mesher ANSA [8]. During the first half of 2023, there was no clear favourite. In its latest release, in July 2023, ANSA has introduced several improvements for high order meshing to the code. The quality of the high-order elements has been improved. Robustness for anisotropic high-order boundary layer elements has been enhanced. ANSA can now automatically split tetrahedral elements into hexahedral elements. It is currently our preferred option.

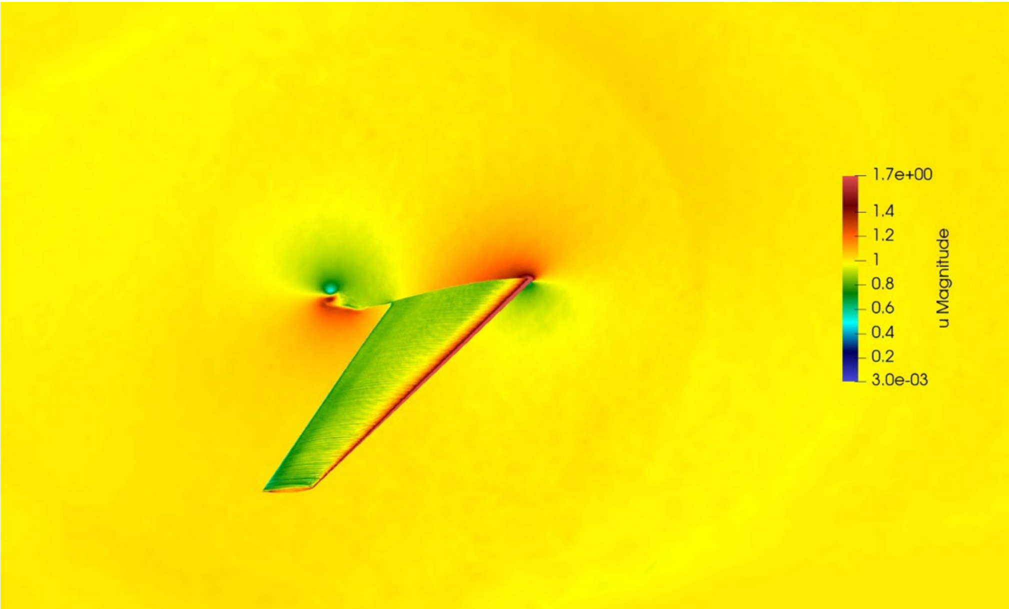

Figure 7: Velocity magnitude over the CRM wing.

Finally, we want to mention that Alya is being interfaced with the Distributed Data-parallel Scientific Visualisation tool VISTLE [7] for in situ visualisation. Further details are reported in D4.1. Moreover, a flexible workflow designed to dynamically manage parallel efficiency during runtime by selecting the appropriate resources based on metrics such as communication efficiency and load balance is being developed in WP3 T3.4 (see D3.1).

5 Next Steps

Our latest improvements make us confident that we should have a robust high-order WMLES workflow soon. We expect to be able to simulate correctly the HLPW Case 1 and at least one or two geometries proposed in Case 2. Due to the computational limitations, we are experiencing with the delay in installing Marenostrum 5, we have recently switched to second-order elements to reduce the number of DoFs in our cases. However, as soon as the situation normalises, we wish to switch back to third or fourth-order elements. It will be interesting to analyse the convergence of the results when the discretisation order is increased. We would also like to obtain mesh converge results for a fixed order and an increasing number of elements.
Following the timeline and Gantt chart we have presented in D2.1, the two critical tasks for the project's second year will be to improve WMLES simulations of more complex geometries and couple the flow simulations with a combustion model. For the second task, the combustion group, led by Daniel Mira, has an open position we hope to fill soon.

Moreover, we will continue working on flow control using classical and AI-based approaches and intelligent postprocessing using POD and DMD.

6 References

[1] C. A. Kennedy and A. Gruber, "Reduced aliasing formulations of the convective terms within the Navier-Stokes equations for a compressible fluid", Journal of Computational Physics, vol. 227, no. 3, pp. 1676–1700, 2008.
[2] J.-L. Guermond, R. Pasquetti, and B. Popov, "Entropy viscosity method for nonlinear conservation laws", Journal of Computational Physics, vol. 230, no. 11, pp. 4248–4267, 2011.
[3] G. Coppola, F. Capuano, S. Pirozzoli, L. de Luca, "Numerically stable formulations of convective terms for turbulent compressible flows", Journal of Computational Physics, Volume 382, 2019, https://doi.org/10.1016/j.jcp.2019.01.007.
[4] "GMSH.", https://gmsh.info
[5] "HOPR.", https://github.com/hopr-framework
[6] "HLPW.", https://hiliftpw.larc.nasa.gov
[7] "VISTLE.", https://vistle.io
[8] "ANSA.", https://www.beta-cae.com/ansa.htm
[9] "AUTOCFD.", https://autocfd.eng.ox.ac.uk
[10] "GeMPa.", https://gitlab.com/dosimont/gempa
[11] Owen, H., Chrysokentis, G., Avila, M., Mira, D., Houzeaux, G., Borrell, R., Cajas, J. C., & Lehmkuhl, O. (2020). Wall-modelled large-eddy simulation in a finite element framework. *International Journal for Numerical Methods in Fluids*, *92*(1), 20-37.https://doi.org/10.1002/fld.4770