HORIZON-EUROHPC-JU-2021-COE-01

EXCELLERAT P2

The European Centre of Excellence for Engineering Applications

Project Number: 101092621

D2.11
Updated 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):	Tufan Arslan	BSC
	Herbert Owen	BSC
Approved by	Executive Centre Management	17.12.2025
Reviewer	Guillaume Daviller	Cerfacs
Reviewer	Fulvio Stella	URMLS
Dissemination Level	Public	

Date	Author	Comments	Version	Status
29.04.2025	Tufan Arslan	First draft	V0.1	Draft
20.05.2025	Herbert Owen	Second draft	V0.2	Draft
02.06.2025	Herbert Owen	Final version	V1.0	Final
16.12.2025	Tufan Arslan	Revised version for resubmission, reviewers' comments have been addressed	V2.0	Final

List of abbreviations

AFC	Active Flow Control
AI	Artificial Intelligence
CAD	Computer Aided Design
CFD	Computational Fluid Dynamics
CO2	Carbon Dioxide
CPU	Central Processing Unit
CRM	Common Research Model
CRM-HL	Common Research Model -High Lift
DMD	Dynamic Mode Decomposition
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
M	Month
NOx	Nitrous Oxide
POD	Proper Orthogonal Decomposition
pyLOM	Python low-order modelling
SPOD	Spectral Proper Orthogonal Decomposition
SVD	Singular Value Decomposition
UC	Use Case
VR	Virtual Reality
WMLES	Wall Modelled Large Eddy Simulation
WP	Work Package

Executive Summary

This document summarises the progress made in the 4th use case (UC-4) Fully integrated aircraft simulations with emissions models from the beginning of EXCELLERAT P2 and at the last section it focuses on the progress made since Month (M) 12.

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 and second 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 had been slightly slower than expected due to a lack of computational resources originated by the delay in the installation of Marenostrum 5. The system is fully operational since March 2024.

The project's second year was 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, opened a position and it has been filled since January 2025. The activity on this task is still going on.

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.

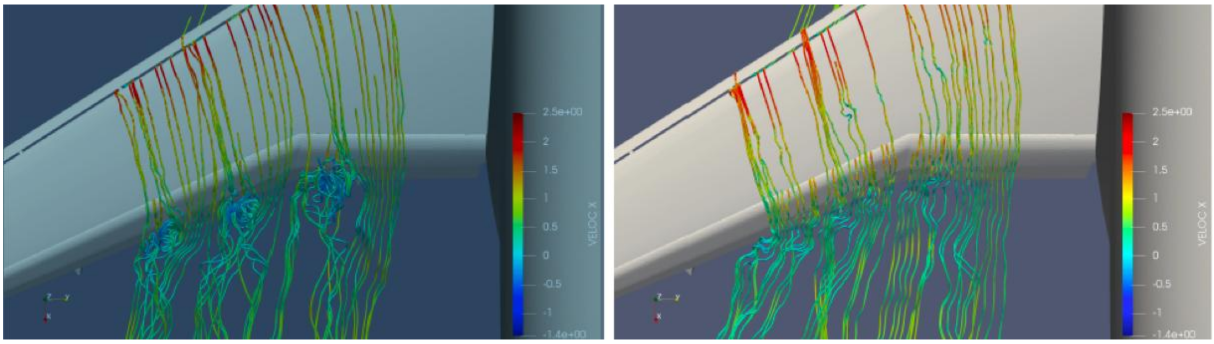

Figure 1: Instantaneous vortical structures calculated with Alya

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 Input Output (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.

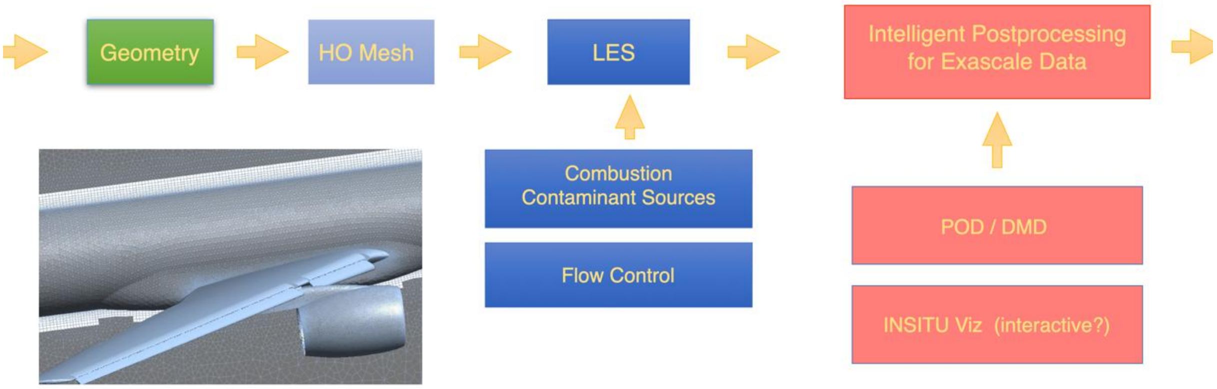

Figure 2: Workflow for fully integrated aircraft simulations with emissions models

The workflow starts with a complex geometry provided in some Computer Aided Design (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. Progress has been made on the coupling between SOD2D and Alya as explained in Section 4, subsection 'Extensions of SOD2D for engine plume characterisation'. A model that deals with multicomponent mixtures has been introduced in SOD2D and the way in which the coupling between both codes will be performed in now clearer.

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], POINTWISE ([12]) 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 since M12

During the M12–M28 period, BSC has worked on three main topics: high order mesh generation, Extensions of SOD2D for engine plume characterisation and development of POD/DMD tool.

Progress in mesh generation

Progress in mesh generation focused primarily on developing high-order meshes for the High Lift Prediction Workshop CRM-HL Case 2.4 (Chord Reynolds Number = 5.4M, Mach = 0.2), with increasing complexity and resolution over time.

Initial work involved generating coarse and medium high-order meshes with 130 million and 760 million degrees-of-freedom (DoF) respectively. A combination of open-source (GMSH) and commercial (ANSA) meshing tools was used. Some details of the highly complex CRM-HL geometry and the resulting mesh can be seen in Figure 3.

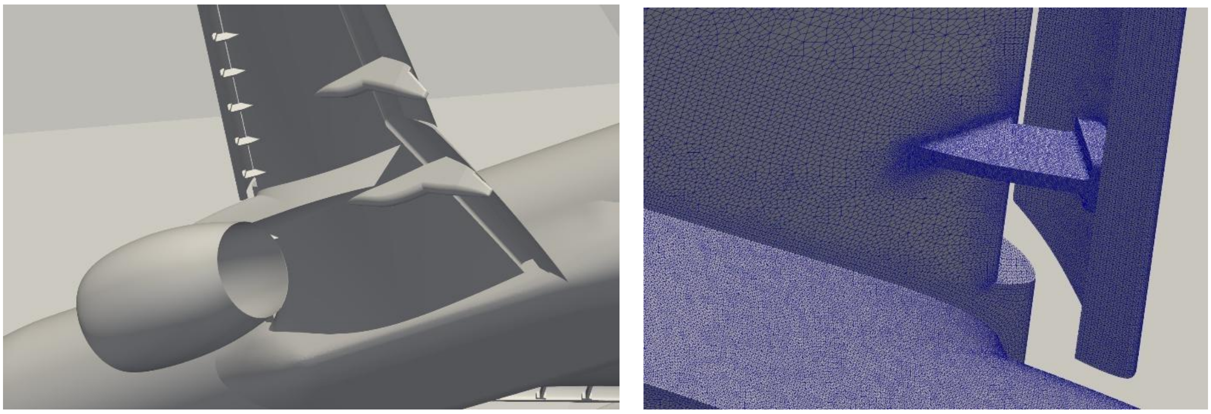

Figure 3: The geometry and the mesh under the wing of CRM-HL Case 2.4 (760M DoF)

Simulations using the quadratic medium mesh with 760 million DoF reached statistical convergence, yielding accurate time-averaged lift and drag coefficients and high-resolution flow visualisations which can be seen in Figure 4.

Figure 4: Flow structures over the surfaces of CRM-HL Case 2.4 (760M DoF)

In the pursuit towards exascale, significantly finer meshes have been constructed to test the limits of commercial meshing software. For highly complex geometries, mesh generation is probably one the main limiters when trying to use a full exascale machine. For the moment the biggest meshes we have constructed are a curved quadratic mesh with 4.9 billion DoF and a 3rd-order high-order mesh with 16 billion DoF. These meshes have been used mainly for scalability tests. We are currently starting a production run with the 4.9 billion DoF mesh, but flow statistics were not yet available at the time of reporting. For even finer meshes, we are working on an in-house tool that can elevate second or third order curved meshes to fourth or higher orders, thus leading to much higher number of DoFs.

Generation of high order meshes for highly complex is an active area of research and tools are far from mature. Significant challenges have been encountered when creating high order meshes near highly curved wall boundaries. In the pursuit of better meshes that can significantly affect the quality of the results, a second commercial meshing software, Pointwise [12], is being evaluated for introduction into our workflow. It has demonstrated strong capabilities in generating curved meshes up to 4th order.

Finally, we would like to mention that we are starting to work on a high order elastic solver within SOD2D, that can be used to improve mesh quality.

Extensions of SOD2D for engine plume characterisation

A new development for extending SOD2D to deal with multicomponent mixtures has been introduced. The code has been tested on canonical problems like the TGV as shown in Figure 5. For this case, two species with the same diffusivities and molecular weight were chosen, so a direct comparison (validation not included here for the sake of brevity) with the single component solution could be made.

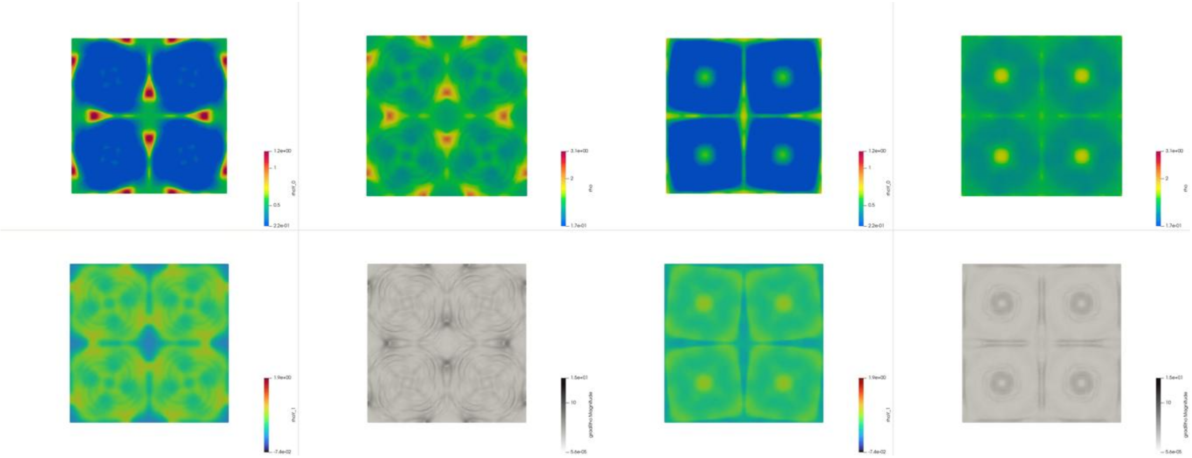

Figure 5: Concentration fields (species Y0 and Y1), velocity gradient and density for a problem setup with size 40x40x40

This enhancement is intended to support the calculation of aircraft engine exhaust plumes with non-reactive and reactive descriptions of the gas composition. A workflow using Alya for the internal flow in the combustion chamber and SOD2D for the plume was proposed, allowing to perform aircraft plume analysis with detailed composition of the engine exhaust byproducts. At this stage, the coupling has not been tested or used in production simulations since the focus was on validation. Currently, conceptual work is being carried out through preliminary simulations using a representative aircraft geometry and relevant operating conditions. These simulations are guided by thermodynamic principles derived from typical turbojet or turbofan engine equations, which are used to estimate outlet conditions at the engine exhaust (see, Figure 6) These outlet conditions are then applied to simplified geometries, and a parametric analysis is conducted to evaluate the impact of various ambient conditions, engine bypass ratios, and angles of attack. Results for this analysis will be reported on the next periodic report.

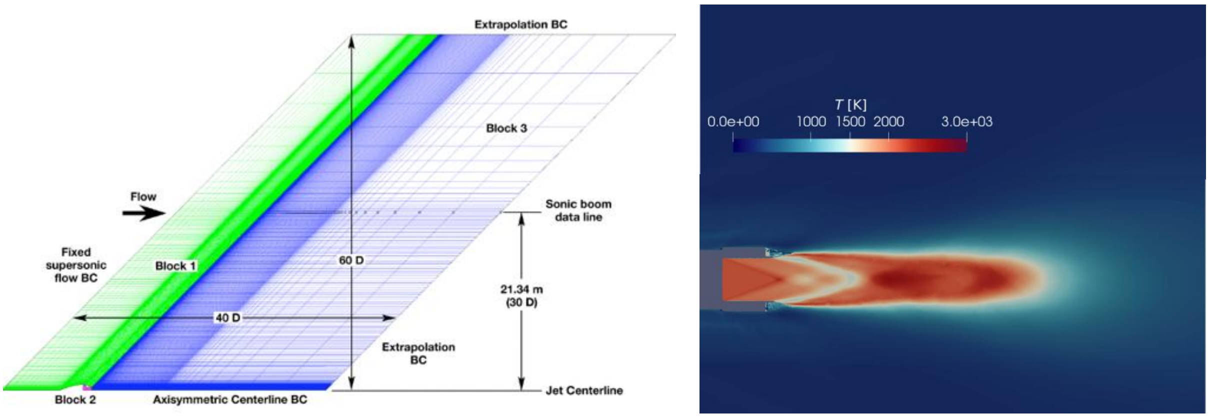

Figure 6: Computational domain and calculated temperature field of simplified nozzle model of the aircraft engine

Development of POD/DMD tool

Runtime post processing can be enhanced by performing modal decompositions of the flow field. To do so, SOD2D can be coupled with Python low-order modelling (pyLOM) [9][13], the reduced order modelling tool developed by BSC in collaboration with the company UPM. pyLOM contains the algorithms of POD, DMD and Spectral Proper Orthogonal Decomposition (SPOD), as well as efficient Singular Value Decomposition (SVD) and matrix-matrix multiplication, all of them tailored for supercomputers. The core of these algorithms is an adaptation of the parallel QR factorisation from Demmel et al. [3] for tall and skinny matrices. A randomised version of this factorisation [14] can be used to reduce the cost of the algorithm and to accumulate the information as the number of snapshots increase. pyLOM can be either installed as a python package to execute its non-compiled version (with GPU support) or can be compiled from the source code to enhance its performance when running on CPU. The in-situ modal decomposition strategy is sending the flow field results every N timesteps from SOD2D to pyLOM through a SmartRedis database. In that way, the SOD2D simulation is executed in the GPU and pyLOM makes use of the idle CPUs. This workflow reduces the amount of disk I/O and allows to build reduced order models with finer time sampling than the a-posteriori approach.

pyLOM has been tested by data created by simple generic fluid problems such 3D unsteady flow around circular cylinder. Figure 12 shows the complex 3D flow structures from various time steps as an input (left) and resulting 2D plots showing the extracted modes from the large instantaneous flow field data.

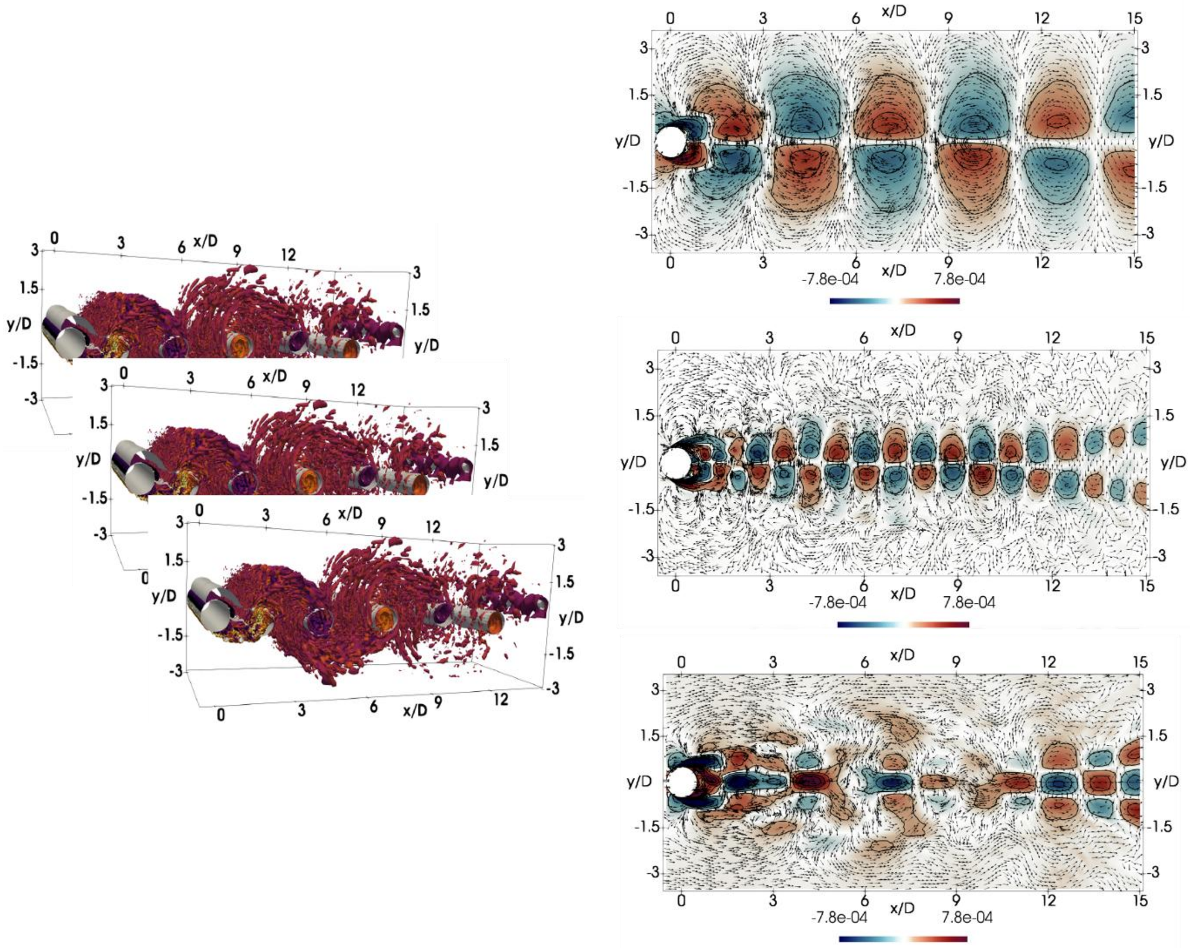

Figure 7: Instantaneous flow field samples (left) and extracted modes (right).

Computing POD and DMD of real industrial cases for aeronautics involves the need of using high-performance computing algorithms and libraries with a good scalability. For this reason, pyLOM interfaces with BLAS and LAPACK for efficient vector and matrix algebra. It also interfaces with the library FFTW and nFFT for efficient Fourier transformation algorithms. When available, pyLOM can interface with Intel MKL for more efficient implementations of these libraries. When deploying into HPC systems, pyLOM can be linked against these libraries already available in the system, hence making available the most efficient algorithms.

Figure 8 (left) shows the speed-up of the SVD for the two versions of pyLOM, the compiled one (which uses C and Cython scripts) and the non-compiled one (based on Python). The test case has been the full-scale LES aircraft, with 70 million grid points and 113 snapshots. It is clearly seen that the code scales correctly up to 4000 cores. Figure 8 (right) shows the speed-up of the whole DMD algorithm. More than 85% of the computation is on the SVD. Hence, the global DMD algorithm scales nearly as well as the SVD algorithm, while other parts of the code are serial and are not shown to be scaling.

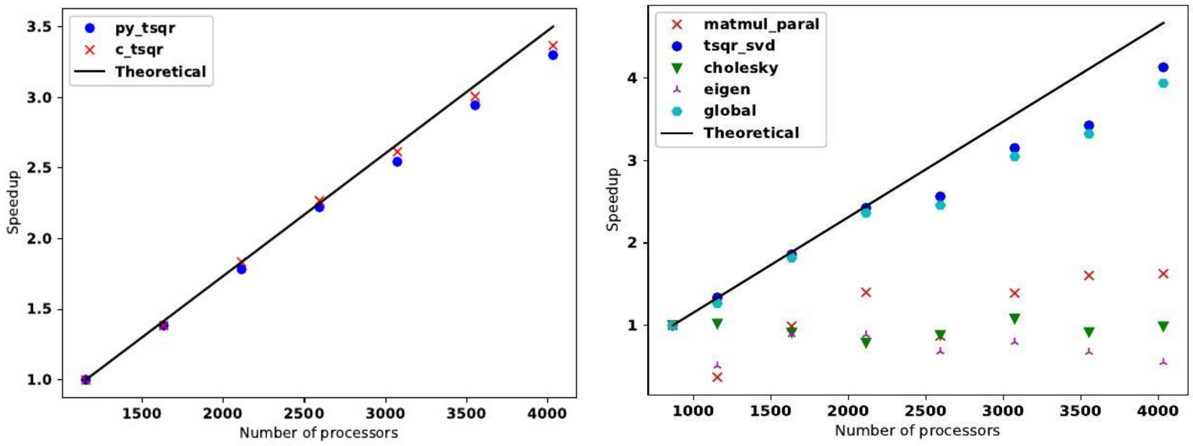

Figure 8: Speed-up plot of the SVD algorithm (left) and DMD algorithm (right) for the full-scale LES aircraft, with 70 million grid points and 113 snapshots.

Finally, these two algorithms were tested on the full-scale LES aircraft case as a demo of the capabilities of pyLOM. These were run with the same number of cores as the simulation (4032) and the SVD took about 3 seconds to complete. The DMD modes obtained are shown in Figure 9. On top there are the real parts of the first three modes of the streamwise velocity while at the bottom there are the imaginary parts.

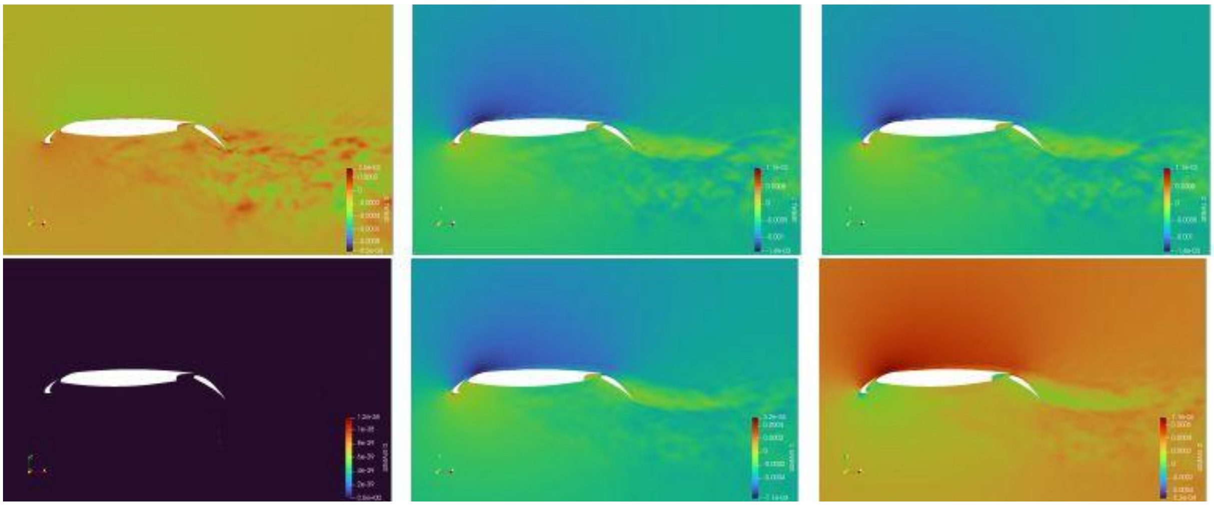

Figure 9: DMD modes of the full-scale LES aircraft, with 70 million grid points and 113 snapshots. On top the real modes and on the bottom their imaginary part. From left to right are modes 1 to 3.

Recent developments and the plans for the part of Active Flow Control (AFC)

The study investigates active flow control (AFC) of a 30P30N high-lift wing at a Reynolds number $Re_c = 450{,}000$ and angle of attack $\alpha = 23^{\circ}$ using wall-resolved large-eddy simulations (LES). Two optimisation strategies will be explored: open loop Bayesian optimisation (BO) and closed-loop deep reinforcement learning (DRL), both targeting the mitigation of stall and the improvement of aerodynamic efficiency via synthetic jets on the slat, main, and flap elements.

The 30P30N high-lift configuration, originally developed by McDonnell Douglas, has become a widely studied benchmark for the aerodynamics and aeroacoustics of multi-element wings. First introduced in the NASA High-Lift Prediction Workshop [15].

The primary objective of the work is to develop and assess an optimisation workflow for AFC using data-driven methods, with a particular focus on Deep Reinforcement Learning (DRL), in order to enhance the aerodynamic performance of the airfoil under high-lift conditions. Initial investigations reveal that the AFC optimisation landscape is highly nonlinear and high-dimensional, involving strongly coupled parameters such as jet location, number of actuators, jet velocity, injection angle, and reward formulation of DRL. To manage this complexity in the early stages of the study, the design space was deliberately reduced by considering two fixed jet locations—one positioned over the slat and one upstream of the flap over the main element—and optimizing only the jet actuation velocity.

The first results demonstrate that BO is an effective optimisation strategy in this constrained setting, achieving measurable aerodynamic improvements with a lift coefficient increase of ΔCl = 0.15% and a drag reduction of ΔCd = −9.69%. Figure 10 shows the alteration in the mean flow field based on BO and the slat wake is eliminated with AFC. These findings provide a reference baseline for subsequent extensions toward higher-dimensional optimisation using DRL-based control strategies.

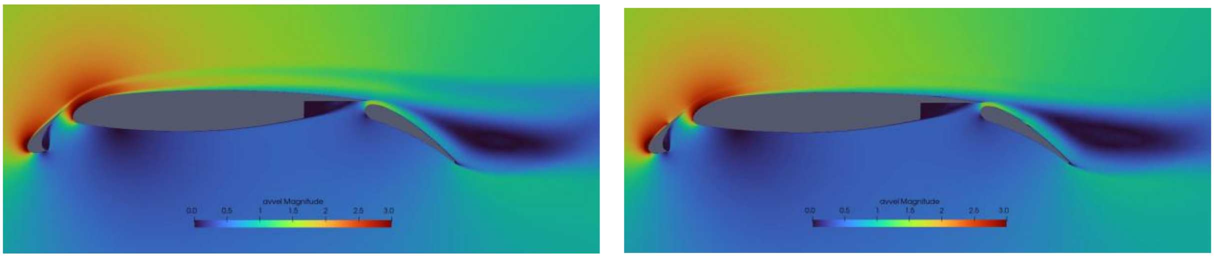

Figure 10: Mean flow field for baseline model (left) and BO model (right).

5 Next Steps

Following the timeline and Gantt chart we have presented in D2.1, the two critical tasks for the project's last year will be to improve WMLES simulations of more complex geometries and couple the flow simulations with a combustion model.

Future efforts will focus on completing the simulations using the 16 billion DoF high-order mesh and extracting detailed flow statistics.

The performance evaluations are planned to scale up to the full Mare Nostrum 5 machine (4096 GPUs). Mesh generation for complex geometries will continue, with ANSA and Pointwise being used in parallel to ensure the creation of high-quality curved high-order meshes. 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.

The species model integrated into SOD2D will be tested, and the coupled Alya–SOD2D system will be employed to compute emissions within the flow field. Moreover, we will continue working on flow control using classical and AI-based approaches and intelligent postprocessing using POD and DMD.

In the last period of the project, For the AFC part of the use-case, we also aim to investigate on training the DRL on a surrogate model and applying the DRL for selecting the best jet locations.

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] Demmel, J., Grigori, L., Hoemmen, M., & Langou, J. (2012). Communication-optimal parallel and sequential QR and LU factorizations. SIAM Journal on Scientific Computing, 34(1), A206-A239.

[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] Eiximeno, B., Miró, A., Begiashvili, B., Valero, E., Rodriguez, I., & Lehmkhul, O. (2025). pyLOM: A HPC open-source reduced order model suite for fluid dynamics applications. *Computer Physics Communications*, *308*, 109459.

[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.

[12] "POINTWISE.",https://www.cadence.com/en_US/home/tools/system-analysis/computational-fluid-dynamics/fidelity.html#fidelity-pointwise

[13] "pyLOM.",https://www.bsc.es/es/research-and-development/software-and-apps/software-list/pyLOM

[14] Erichson, N. B., Voronin, S., Brunton, S. L., & Kutz, J. N. (2019). Randomized matrix decompositions using R. *Journal of Statistical Software*, *89*, 1-48.

[15] Klausmeyer S M and Lin J C 1997 Comparative results from a CFD challenge over a 2D threeelement high-lift airfoil Tech. Rep. NASA TM-112858 NASA Langley Research Center NASA Technical Memorandum.