HORIZON-EUROHPC-JU-2021-COE-01

EXCELLERAT P2

The European Centre of Excellence for Engineering Applications

Project Number: 101092621

D2.2 Report on Use-Case Execution

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:	WP2	Use-Case Execution
Author(s):	Matthias Meinke, Ansgar Niemöller	RWTH
	Adam Peplinski, Ronith Stanly, Eman Bagheri, Adalberto Perez, Philipp Schlatter	KTH
	Antoine Dauptain, Mohamed Ghenai	CERFACS
	Michael Wagner	DLR
	Tufan Arslan, Herbert Owen	BSC
	Matic Brank	UL
	Francesco Salvadore, Sergio Pirozzoli, Giulio Soldati	CINECA
Approved by	Executive Centre Management	30-06-2025
Reviewer 1	Guillaume Daviller	CERFACS
Reviewer 2	Simone Bnà	CINECA
Dissemination Level	Public	

Date	**Author**	**Comments**	**Version**	**Status**
30-05-2025	Matthias Meinke, Ansgar Niemöller	Initial draft for internal review	V0.1	Draft
27-06-2025	Matthias Meinke, Ansgar Niemöller	Revised version, including the reviewers' feedback	V0.2	Draft
30-06-2025	Matthias Meinke, Ansgar Niemöller	Final version for submission	V1.0	Final

List of abbreviations

ACARE	Advisory Council for Aeronautical Research in Europe
ACF	Autocorrelation functions
AMR	Adaptive mesh refinement
APU	Accelarated processing unit
ASMR	Automatic static mesh refinement
BSC	Barcelona Supercomputing Center
CAA	Computational aeroacoustic
CCC	Center coordination committee
CCRT	Centre de calcul recherche et technologie
CEA	Commissariat à l’énergie atomique et aux énergies alternatives
CERFACS	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique
CFD	Computational fluid dynamics
CPU	Central processing unit
DLR	German Aerospace Center
DNS	Direct numerical simulation
DRL	Deep reinforcement learning
DMD	Dynamic mode decomposition
DOF	Degree of freedom
ECFD7	7th Extreme CFD Workshop, 2024
FA	Further application
FASTER	H2020 project: First responder advanced technologies for safe and efficient emergency response
FhG	Fraunhofer-Gesellschaft
FPGA	Field Programmable Gate Array
FSDM	FlowSimulator DataManager
FWH	Ffowcs Williams-Hawkings
GPU	Graphics processing unit
HPC	High-performance computing
HIP	A C++ runtime API and kernel language
ITER	International Thermonuclear Experimental Reactor
JSON	JavaScript object notation
KTH	Royal Institute of Technology
LES	Large-eddy simulation
m-AIA	multi-physics PDE solver framework with a focus on fluid dynamics equations
MDA	Multi-disciplinary analysis
MOS	Method of snapshots
MPI	Message passing interface
NACA	National Advisory Committee for Aeronautics
NOx	Nitrogen oxides
OpenACC	Organization for Open Accelerated Computing
OpenMPI	An open source message passing interface
POD	Proper orthogonal decomposition
RANS	Reynolds-averaged Navier-Stokes
RWTH	RWTH Aachen University
SAneg	Spalart-Allmaras one-equation turbulence model in its negative form
SEM	Spectral element method
SIMD	Single instruction, multiple data architecture
SOD2D	Spectral high-order code to solve partial differential equations

StTW	Streamwise-traveling waves
SVD	Singular value decomposition
SPL	Sound pressure level
TOML	A file format
TOPAZE	Supercomputer at the CCRT
UAV	Unmanned aerial vehicle
UC	Use case
UL	University of Ljubljana
UQ	Uncertainty quantification
WEST	Tokamak fusion reactor
VR	Virtual reality
VTK	Visualization toolkit

Executive Summary

This document summarises the achievements of the use case execution up to M30 in the EXCELLERAT P2 project. In total, six use cases and one further application are defined in the proposal. A detailed description of the work on each of the use cases and the further application is provided in dedicate deliverables: D2.4, 2.6, 2.8, 2.10, 2.12, 2.14, 2.16 for reporting period 1 (M1-M12) and D2.5, 2.7, 2.9, 2.11, 2.13, 2.15, and 2.17 for reporting period 2 (M13-M30). Therefore, this document provides a higher-level overview of the individual use cases, for more detailed technical information the reader is referred to the aforementioned deliverables.

In general, all use cases are on track, at present there are no major problems or delays in the use cases, which would pose a risk for a successful completion within the project. Minor amendments have been done in terms of codes being replaced and timelines which were adapted.

Table of Contents

List of Figures

List of Tables

1 Introduction

The document summarizes the workflow developments for all use cases (UCs) in the EXCELLERAT P2 project. Since the progress achieved in each use case up to M30 is described in separate deliverables, this document will not contain any technical detail to avoid a duplication of contents across the deliverables within EXCELLERAT P2. The main idea of this document is to provide a concise overview of the current state in the workflow development highlighting the achievements, impact and the plans for the next developments.

Table 1 presents an overview of the six use cases and one further application with the respective topic addressed, responsible partner, as well as the used code.

	Topic	**Partner**	**Code**
UC-1	External aircraft aerodynamics	DLR	CODA
UC-2	Hydrogen combustion for propulsion	CERFACS	AVBP
UC-3	Mitigation of aeroacoustic noise	RWTH	m-AIA
UC-4	Fully integrated aircraft simulations with emission models	BSC	Alya
UC-5	High-fidelity simulations of rotating parts	KTH	Neko
UC-6	Active control for drag reduction of transonic airfoils	CINECA	STREAmS
FA-1	Engineering design and digital twin of the first wall of a tokamak fusion reactor	UL	L2G, OpenFoam, Raysect

Table 1: Use cases and further application showing the code owners and used simulations codes

2 UC-1 External Aircraft Aerodynamics

Partner: DLR, Code: CODA

2.1 Overview

The EXCELLERAT P2 project represents a critical advancement in European high-performance computing capabilities for aerospace applications. During the second reporting period, the CODA Application Use Case (UC-1) has demonstrated remarkable progress in developing next-generation Computational Fluid Dynamics (CFD) solutions for aircraft design and aerodynamic simulations.

CODA (CFD by Onera, DLR, and Airbus) is positioned as the successor to the TAU CFD package, which has been the industry standard for over 20 years and was instrumental in designing aircraft such as the Airbus A380 and A350. The project focuses on enhancing CODA's scalability and performance to meet the demands of pre-exascale and exascale computing environments, enabling more accurate and efficient aircraft design processes.

The use case simulates steady airflow at transonic speeds around aircraft geometries, solving Reynolds-averaged Navier-Stokes (RANS) equations with advanced turbulence modeling. Test cases range from tens of millions to several billion mesh elements, representing both research-scale and production-scale computational challenges. The workflow integrates three core components: the CODA CFD solver, the FlowSimulator framework for multi-disciplinary analysis, and the Spliss sparse linear system solver.

This comprehensive approach addresses the critical need for high-precision CFD simulations in modern aircraft development, where computational efficiency directly impacts design timelines, cost reduction, and the ability to explore innovative aerospace technologies. The project's emphasis on scalability and performance optimization positions European aerospace research and industry at the forefront of computational aerodynamics capabilities.

2.2 Key Achievements

Scalability Breakthroughs

The most significant achievement has been the dramatic improvement in CODA's parallel scalability. Strong scaling tests demonstrated 83% parallel efficiency on 32,768 cores, representing a 22% improvement over baseline performance (61% efficiency). Weak scaling results were even more impressive, achieving 96% parallel efficiency on the same scale, marking a 24% improvement from the 72% baseline. These results were validated on DLR's CARA HPC system using AMD Naples architecture, confirming CODA's readiness for pre-exascale computing environments.

Mixed-Precision Computing Innovation

The implementation of mixed-precision arithmetic in the Spliss linear solver has yielded exceptional performance gains. By utilizing single-precision calculations for inner loops (Jacobi preconditioner) while maintaining double-precision for outer loops (GMRES), the system achieves up to 72% acceleration in calculation time. This optimization leverages doubled SIMD instruction efficiency and halved memory bandwidth requirements, with users reporting average speedups of approximately 30% for complete simulations.

Containerized Deployment Success

The development of a comprehensive containerized workflow represents a major advancement in software deployment and portability. The entire CODA, FlowSimulator, and Spliss ecosystem can now be delivered as a single container image, enabling system-independent execution across diverse HPC platforms. Performance benchmarks confirm that containerized deployments maintain comparable efficiency to native installations while dramatically improving deployment flexibility and user accessibility.

Advanced Mesh Partitioning Capabilities

The implementation of hierarchical mesh partitioning has revolutionized large-scale simulation capabilities. The three-level distribution approach (across nodes, MPI processes, and threads) delivers up to one order of magnitude speedup in mesh partitioning operations. This breakthrough enables simulations of meshes exceeding one billion elements and successful execution on 131,072 cores, representing the full capacity of DLR's CARO system without degrading load balance or solver performance.

Multi-Architecture Validation

Comprehensive performance evaluation across emerging CPU and GPU architectures ensures CODA's adaptability to future computing environments. Testing encompasses AMD Zen series, Intel Sapphire Rapids, ARM-based processors (Graviton 2-4, Nvidia Grace), and GPU accelerators (Nvidia H100/A100, AMD Mi210). This extensive validation provides critical insights for optimization strategies and future HPC system design decisions.

2.3 Impact and Significance

Technical Impact

The achieved improvements fundamentally transform CODA's computational capabilities and position it as a leading-edge CFD solution for exascale computing. The 22-24% scalability improvements directly translate to reduced computational time and energy consumption for large-scale aerodynamic simulations. The mixed-precision computing breakthrough offers immediate productivity benefits, with 30% average speedups significantly accelerating design iteration cycles in aerospace development.

The billion-element mesh simulation capability represents a quantum leap in computational scale, enabling previously impossible high-fidelity simulations of complex aircraft configurations. This capability supports the analysis of complete aircraft systems with unprecedented detail, facilitating breakthrough innovations in aerodynamic design and performance optimization.

Industrial Significance

For the European aerospace industry, these achievements provide substantial competitive advantages. Reduced simulation times enable more extensive design space exploration, leading to more fuel-efficient, quieter, and safer aircraft designs. The improved scalability ensures that European research organizations and industry partners can fully utilize current and future HPC investments, maximizing return on computational infrastructure.

The containerized deployment capability addresses a critical industry need for consistent, portable simulation environments across diverse computing platforms. This innovation reduces deployment complexity, accelerates user onboarding, and ensures consistent results across different computing environments, from local workstations to large-scale HPC systems.

Strategic Significance

The project's success reinforces Europe's leadership in computational aerodynamics and high-performance computing applications. By demonstrating exascale readiness, CODA positions European aerospace capabilities at the forefront of global competition. The multi-architecture validation ensures adaptability to emerging computing technologies, providing strategic flexibility for future technological evolution.

The integration of advanced computing techniques (mixed-precision, hierarchical partitioning, containerization) establishes best practices that can be applied across the broader HPC and CFD communities. These innovations contribute to the European Centre of Excellence mission by advancing both specific application capabilities and general HPC methodologies.

2.4 Next Developments

Immediate Priorities

The project will continue systematic evaluation of CODA and FlowSimulator efficiency improvements, implementing scalability regression testing for all new releases to maintain performance standards. Containerized workflow deployment will be extended to additional HPC systems, validating performance consistency across diverse computing environments and expanding accessibility for European research institutions and industry partners.

Advanced Architecture Exploration

Ongoing evaluation of emerging CPU and GPU architectures will continue, focusing on early-access systems that represent the next generation of HPC capabilities. This proactive approach ensures CODA's optimization for future computing platforms before they become widely available, maintaining European competitive advantages in computational aerodynamics.

Scalability Demonstrator Development

The project will develop a comprehensive scalability and efficiency demonstrator for CODA and FlowSimulator, showcasing capabilities on the largest available systems. This demonstrator will serve as a validation platform for exascale readiness and provide a benchmark for performance evaluation across different computing environments.

Workflow Enhancement

Continued development of FlowSimulator components will focus on supporting increasingly large meshes and core counts, with particular emphasis on improving memory efficiency and communication optimization. These enhancements will ensure robust performance as simulations scale to exascale levels.

Production Readiness

The transition from the research prototype to a production-ready system will involve extensive validation with industrial partners, ensuring that performance improvements translate effectively to real-world aerospace design workflows. This includes integration with existing design processes and validation of results against established benchmarks and experimental data.

The successful completion of these developments will establish CODA as the premier European CFD solution for next-generation aerospace design, supporting the industry's transition to more efficient, sustainable, and innovative aircraft technologies while maintaining Europe's leadership in computational aerodynamics and high-performance computing applications.

3 UC-2 Hydrogen combustion

Partner: CERFACS, Code: AVBP

3.1 *Overview*

The EXCELLERAT P2 AVBP Use Case (UC-2) addresses one of the most critical challenges facing the European aerospace and energy industries: the transition from fossil fuels to hydrogen in gas turbine technology. This transition is essential for achieving EU decarbonization objectives within the next 20 years, but it presents unprecedented technical challenges due to hydrogen's unique combustion characteristics - burning faster, diffusing faster, and exploding more readily than conventional fuels.

The use case focuses on developing and demonstrating an integrated workflow for hydrogen combustion simulations using the AVBP solver, specifically targeting exascale High-Performance Computing (HPC) architectures. AVBP, developed at CERFACS, is a leading Computational Fluid Dynamics (CFD) solver specialized in Large Eddy Simulation (LES) of complex combustion phenomena. The project aims to enable realistic industrial-grade simulations of hydrogen combustion chambers on pre-exascale and exascale EuroHPC systems.

The technical approach centres on Automated Static Mesh Refinement (ASMR) workflows, which represent a paradigm shift from traditional manual mesh generation to intelligent, simulation-driven mesh adaptation. This approach is critical because the quality of computational grids fundamentally determines the accuracy and reliability of combustion simulations. The workflow integrates three core components: the AVBP CFD solver for combustion simulation, the KalpaTARU library for massively parallel mesh adaptation, and the Lemmings workflow automation system for seamless orchestration of complex computational processes.

The strategic importance of this work cannot be overstated. European companies face extremely tight schedules for redesigning combustion chambers for hydrogen applications, and computational simulation will be the primary tool for achieving these design objectives. Success in this use case will provide European industry with world-leading capabilities for hydrogen combustion simulation, supporting the continent's transition to sustainable energy systems while maintaining industrial competitiveness.

3.2 *Key Achievements*

Automated Static Mesh Refinement Workflow Development

The most significant achievement has been the successful development and validation of the ASMR workflow using the Lemmings automation framework. This workflow represents a fundamental advancement in computational combustion methodology, automatically transitioning between simulation and mesh adaptation steps based on statistical convergence of flow features. The iterative process begins with a coarse initial mesh, performs LES simulation until convergence, then adapts the mesh to better capture quantities of interest, creating progressively refined computational grids.

Multi-Scale Validation Success

The workflow has been comprehensively validated across multiple complexity levels, demonstrating both technical robustness and practical applicability. Initial validation used the TRAPVORTEX simplified geometry, containing all conditions of real industrial LES (complex boundary conditions, jet cross-flow injection, annular geometry) within a manageable 200,000 tetrahedral elements. This validation successfully converged to a 2.5 million element mesh with adequate flow feature representation, providing rapid prototyping and validation capabilities.

Large-Scale Hydrogen Combustion Demonstration

Advanced validation was achieved using the TUB technically premixed H2/air injector, a laboratory-scale hydrogen combustor from Technische Universität Berlin. The ASMR workflow demonstrated exceptional efficiency, progressively adapting from an initial 2.8 million cell mesh through 5.7M, 11.5M, 23M, and finally 46 million elements. Critically, this automated approach achieved comparable accuracy to expert-generated 200 million element meshes while reducing computational cost by a factor of five, representing a transformative improvement in simulation efficiency.

Multi-Platform HPC Integration

The workflow has been successfully extended across multiple European HPC systems, demonstrating portability and scalability essential for exascale deployment. Integration has been achieved on TOPAZE cluster (CCRT), LUMI-C, and LEONARDO (CINECA), with systematic testing validating performance consistency across diverse architectures. This multi-platform capability ensures broad accessibility for European research institutions and industrial partners.

Advanced Workflow Automation Enhancements

Significant improvements have been implemented in the Lemmings workflow automation system, including transition from ParMETIS to PT-Scotch graph partitioner (resolving licensing concerns), development of generic YAML-driven configuration scripts, and automatic computation of refinement steps based on initial and target meshes. These enhancements eliminate manual intervention requirements and improve workflow reliability and reproducibility.

KalpaTARU Parallel Mesh Adaptation Progress

The transition from the sequential mesh adaptation tool PyHIP to massively parallel KalpaTARU has been initiated, addressing the critical need for exascale-compatible mesh manipulation capabilities. Development includes resolution of partitioning conflicts, support for periodic adaptation, and prototype parallel job integration. This advancement is essential for handling meshes exceeding 200 million elements required for exascale demonstrations.

3.3 Impact and Significance

Industrial Transformation Potential

The ASMR workflow achievements represent a paradigm shift for European combustion system design, particularly for hydrogen applications. The demonstrated five-fold reduction in computational cost while maintaining simulation accuracy directly addresses industry needs for rapid design iteration under tight development schedules. This efficiency improvement enables more extensive design space exploration, accelerating the development of hydrogen-compatible combustion systems essential for decarbonization objectives.

Technical Leadership in Hydrogen Combustion

The successful validation of hydrogen combustion simulations using automated mesh refinement establishes European leadership in this critical technology domain. The ability to accurately simulate hydrogen's unique combustion characteristics (faster burning, diffusion, and explosion) provides European aerospace and energy companies with competitive advantages in developing next-generation propulsion and power generation systems.

Exascale Computing Readiness

The multi-platform validation and parallel algorithm development demonstrate clear progress toward exascale computing capabilities. The workflow's successful operation on leading European HPC systems (LUMI, LEONARDO) validates the technical approach and confirms readiness for deployment on future exascale platforms. This readiness is crucial for maintaining European competitiveness in high-performance computational combustion.

Workflow Automation Innovation

The Lemmings automation framework represents a significant advancement in computational workflow management, with applications extending beyond combustion simulation to broader CFD and engineering simulation domains. The intelligent automation of simulation-adaptation cycles eliminates traditional bottlenecks in computational design processes, enabling more efficient utilization of HPC resources.

Collaborative Impact and Knowledge Transfer

Positive feedback from Safran Group end-users and successful deployment during the ECFD7 workshop demonstrate effective knowledge transfer from research to industrial application. The systematic collection of job data points and improved benchmark exploitability across different compilation setups enhance the broader combustion simulation community's capabilities.

Strategic European Advantage

The project's success reinforces Europe's position as a leader in both hydrogen technology development and high-performance computing applications. By providing European industry with world-class simulation capabilities for hydrogen combustion, the project directly supports EU strategic objectives for energy transition and technological sovereignty.

3.4 Next Developments

Exascale Mesh Adaptation Scaling

The immediate priority involves completing the transition from sequential HIP to massively parallel KalpaTARU for handling meshes exceeding 200 million elements. This transition is critical for exascale compatibility and will enable simulation of full-scale industrial combustion systems with unprecedented detail. Development will focus on optimizing parallel efficiency and ensuring robust performance across diverse HPC architectures.

GPU Acceleration and Portability Enhancement

Significant effort will be directed toward strengthening GPU robustness and portability, with AMD GPU full support representing a major milestone. This development addresses the reality that future exascale systems will rely heavily on GPU acceleration, and ensuring broad GPU compatibility is essential for workflow deployment across European HPC infrastructure.

Multi-System Exascale Demonstration

The workflow will be demonstrated across multiple EuroHPC systems, validating performance consistency and identifying system-specific optimization opportunities. This comprehensive validation will establish best practices for exascale deployment and provide guidance for future system procurement and configuration decisions.

Industrial Integration and Real Geometry Implementation

Following the M24 milestone, the workflow will be extended to handle real industrial geometries provided by industrial partners, moving beyond simplified validation cases to production-relevant applications. This transition will require addressing confidentiality concerns while maintaining the workflow's generality and reusability for diverse hydrogen combustion applications.

Advanced Algorithmic Development

Continued development of mesh adaptation algorithms will focus on improving periodic adaptation capabilities, enhancing load balancing for complex geometries, and optimizing memory utilization for large-scale simulations. These improvements will ensure robust performance as simulations scale to billions of computational elements.

Comprehensive Benchmarking and Validation

Systematic collection of performance data across different compilation setups, hardware configurations, and problem scales will establish comprehensive benchmarking capabilities. This data will inform optimization strategies and provide quantitative guidance for users selecting appropriate computational resources for specific simulation requirements.

Workflow Ecosystem Expansion

Development will extend beyond core simulation capabilities to include enhanced pre- and post-processing tools, improved integration with design optimization frameworks, and enhanced coupling with experimental validation data. These extensions will create a comprehensive simulation ecosystem supporting the entire hydrogen combustion system development process.

The successful completion of these developments will establish the AVBP ASMR workflow as the premier European solution for hydrogen combustion simulation, directly supporting industry needs for rapid, accurate, and cost-effective design of next-generation combustion systems essential for achieving European decarbonization objectives.

4 UC-3 Mitigation of Aeroacoustic Noise

Partner: RWTH, Code: m-AIA

4.1 Overview

The EXCELLERAT P2 project's m-AIA (multi-physics AIA) application use case represents a cutting-edge initiative in computational aeroacoustics, focusing on the optimization of chevron nozzles for aircraft jet noise reduction. This use case addresses one of the aviation industry's most pressing challenges: achieving the 65% noise reduction target established by the Advisory Council for Aeronautical Research in Europe (ACARE) in their Flightpath 2050 vision.
The project employs a sophisticated multiphysics approach that couples computational fluid dynamics (CFD) with computational aeroacoustics (CAA) to predict jet flow and acoustic far-field noise spectra accurately. The simulation framework tackles the complex challenge of chevron nozzle shape optimization under constraints, where noise reduction must be balanced against thrust loss considerations. Operating at Mach 0.9 and Reynolds number of 1 million, the simulations represent almost realistic flight conditions for commercial aircraft.
Central to this effort is the development of an AI-based optimization workflow that can efficiently navigate the expensive objective function evaluations required for aeroacoustic predictions. The multiphysics simulation framework m-AIA serves as the computational backbone, enabling large-scale simulations on pre-exascale high-performance computing (HPC) systems. The project represents a significant step toward making shape optimization feasible for noise-critical applications that traditionally required prohibitively expensive computational resources.
The collaborative effort between RWTH Aachen University and Fraunhofer Institute (FhG) demonstrates the integration of advanced computational methods with artificial intelligence techniques. This combination addresses the fundamental challenge that noise reduction optimization has remained largely unsolved due to the computational complexity of accurately predicting sound pressure levels in the far field through turbulence-scale resolving simulations.

4.2 Key Achievements

Enhanced Large-Scale Computing Capabilities

The project has achieved remarkable scalability in computational aeroacoustics, successfully executing simulations with up to 3.7 billion CFD cells and 4.9 billion degrees of freedom for CAA calculations. Most significantly, these simulations demonstrated linear scalability on the Hawk HPC system using up to 4,096 nodes totaling 524,000 CPU cores. This represents one of the largest coupled CFD-CAA simulations ever performed and proves that full HPC systems can be effectively utilized for aeroacoustic predictions at pre-exascale levels.

Enhanced Acoustic Prediction Accuracy

Major improvements were made to the Ffowcs Williams-Hawkings (FWH) solver, including the identification and correction of critical formulation errors that previously led to erroneous results. The implementation of Welch's method for power spectral density estimation has significantly improved far-field noise prediction capabilities. Validation studies demonstrated strong agreement with experimental reference data across wide frequency ranges, with results showing insensitivity to FWH surface parameters, confirming the robustness of the approach.

Workflow Automation

A comprehensive Python-based automation system has been developed for the entire m-AIA simulation pipeline. This system manages complex workflows involving multiple simulation phases, various output files, and large data volumes without requiring user supervision. The modular design using JSON workflow definitions and TOML configuration files enables seamless integration into optimization loops, making previously manual and error-prone processes fully automated and reliable.

Successful Hardware Architecture Adaptation

The project achieved major progress in code optimization through successful porting to GPU/APU architectures, specifically AMD MI300A systems. The restructuring of compute kernels to loop over individual nodes rather than full elements resulted in a remarkable 6x performance improvement while maintaining 70% parallel efficiency across 64 nodes and 256 APUs on the Hunter HPC system. This advancement ensures the framework's compatibility with future exascale computing architectures.

AI-Driven Optimization Framework Development

A sophisticated optimization workflow has been implemented using Bayesian optimization with Gaussian Process Regression, featuring fully automated and parallelized execution capabilities. The modular design supports simultaneous evaluation of multiple candidate solutions, dramatically improving optimization efficiency. Testing on NACA airfoil cases has validated the approach, with the framework successfully handling complex constraints and demonstrating the ability to balance exploration and exploitation in the optimization process.

4.3 Impact and Significance

Transforming Aerospace Engineering Design

This work represents a paradigm shift in aerospace engineering, making previously impossible design optimizations feasible for noise-critical applications. The ability to perform accurate aeroacoustic predictions at the scale demonstrated opens new possibilities for aircraft engine design, directly contributing to the aviation industry's sustainability goals. The 65% noise reduction target established by ACARE becomes achievable through the systematic optimization capabilities now available.

Advancing High-Performance Computing Applications

The demonstrated scalability to 4,096 nodes with linear performance characteristics establishes new benchmarks for multiphysics simulations in the HPC community. The successful coupling of CFD and CAA solvers at this scale proves that complex, tightly-coupled multiphysics problems can effectively utilize the full capacity of modern supercomputing facilities. This achievement paves the way for other computationally intensive engineering applications to leverage similar approaches.

Pioneering AI Integration in Engineering Simulation

The successful integration of AI-based optimization with high-fidelity physics simulations represents a significant advancement in computational engineering. The Bayesian optimization framework with parallel evaluation capabilities demonstrates how machine learning can make traditionally intractable optimization problems solvable within reasonable timeframes. This approach serves as a model for other industries facing similar computational challenges.

Enabling Exascale Readiness

The GPU/APU porting achievements position the framework for future exascale computing systems. The 6x performance improvement and demonstrated scalability across different architectures ensure that the computational methods developed will remain relevant and efficient as HPC technology continues to evolve. This forward-looking approach guarantees long-term viability and continued impact.

Establishing New Standards for Multiphysics Simulation

The validated coupling methodology between CFD and CAA solvers, combined with improved FWH acoustic prediction methods, establishes new standards for accuracy and reliability in computational aeroacoustics. The demonstrated agreement with experimental data across wide frequency ranges provides confidence for industrial applications and regulatory acceptance.

4.4 Next Developments

Comprehensive Baseline Comparisons and Analysis

The immediate focus involves conducting large-scale aeroacoustic simulations for the SMC000 round baseline nozzle to establish comprehensive comparative studies against the SMC001 chevron nozzle results. This comparative analysis will provide crucial insights into noise source mechanisms and validate the optimization framework's effectiveness. The data generated will serve as the foundation for understanding the physical mechanisms underlying jet noise generation and reduction.

Completing GPU/APU Architecture Optimization

Ongoing development will focus on completing the GPU/APU porting of the coupled finite volume and discontinuous Galerkin solvers. Based on the initial success with the MI300A architecture, further performance optimization will be implemented to achieve full exascale readiness. These improvements will ensure maximum efficiency across diverse computing architectures and enable broader accessibility of the simulation capabilities.

Industrial-Scale Optimization Implementation

The project will transition from proof-of-concept to industrial-scale implementation by performing initial chevron nozzle optimization runs using the developed AI framework. This phase will involve close collaboration between RWTH and FhG to refine the optimization algorithms based on real-world constraints and requirements. The focus will be on grey-boxing approaches and established methods with proven performance records in similar optimization problems.

Advanced Constraint Handling and Robustness

Future development will enhance the optimization framework's ability to handle both known constraints (such as geometric limitations) and unknown constraints (such as infeasible configurations). This includes improving the robustness of the Bayesian optimization approach and implementing advanced acquisition functions that better balance exploration and exploitation for complex engineering design spaces.

Scaling Toward Exascale Computing

The ultimate goal involves preparing for true exascale computing capabilities through continued algorithm refinement and hardware adaptation. This includes optimizing data management strategies for the massive datasets generated, improving fault tolerance for long-running optimizations, and developing more sophisticated load balancing approaches for heterogeneous computing environments. The project aims to demonstrate that complex multiphysics optimization problems can be solved efficiently using the next generation of supercomputing systems.

4.5 Updated Gantt Chart

The timeline of the developments in UC3 has been slightly shifted. While parts of the optimization workflow developments have been shifted to a later time, large scale simulation runs have been performed earlier due to the fact that the large scale HPC system HAWK was taken out of operation in April 2025. The updates are reflected in the Gantt Chart in Figure 1.

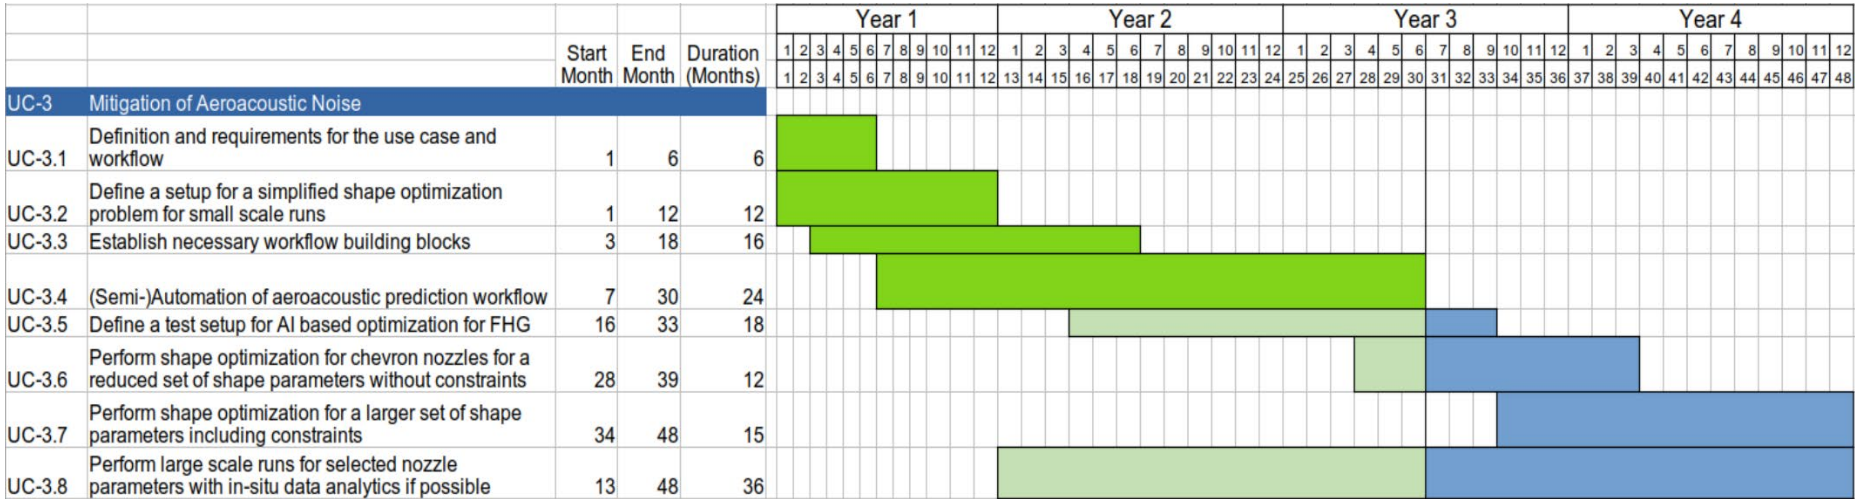

Figure 1: Updated Gantt Chart for UC3

5 UC-4 Fully integrated aircraft simulations with emissions models

Partner: BSC, Code: Alya

5.1 Overview

The EXCELLERAT P2 project represents a significant advancement in computational fluid dynamics (CFD) for aerospace applications, specifically focusing on Use Case 4: "Fully integrated aircraft simulations with emissions models." This initiative directly addresses the European Union's decarbonisation objectives for air transportation by 2050, as outlined in Flightpath 2050, which mandates substantial reductions in CO2, NOx emissions, and aircraft noise.

The project centres on developing a comprehensive multiphysics platform that integrates advanced simulation techniques using leading-edge High Performance Computing (HPC) architectures. The core objective is to enable digital twins of fully integrated aircraft simulations by coupling external aerodynamics with engine emissions simulations through scale-resolving simulation methods. This includes incorporating flow control devices, high-fidelity combustion simulations, and deep reinforcement learning for flow control of high-lift wings. Central to this effort is the development of a new high-order version of the Alya CFD code, transitioning from the standard low-order finite element version used in EXCELLERAT P1. The new implementation utilizes a spectral element method of arbitrary order with Lobatto-Gauss-Legendre quadrature, incorporating advanced stabilization techniques and GPU optimization through OpenACC. This evolution represents a fundamental shift toward more accurate, energy-efficient computational methods capable of handling the complexity required for next-generation aircraft design.

The workflow encompasses the entire simulation pipeline, from complex CAD geometry processing through mesh generation, Wall-Modelled Large Eddy Simulation (WMLES), combustion modelling, flow control implementation, and intelligent post-processing using advanced data analytics tools including Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD).

5.2 Key Achievements

High-Order Method Development and Validation

The project successfully developed and validated a new high-order spectral element version of Alya, demonstrating superior accuracy compared to its low-order predecessor. Comprehensive validation was conducted across problems of increasing complexity, including the Taylor Green vortex at multiple Mach numbers (Ma=0.1 and Ma=1.25), channel flow at Reynolds number Re_tau=950, turbulent flow behind cylinders, and complex engineering applications such as the Windsor model car simulation with 150 million degrees of freedom.

Large-Scale Computational Capabilities

Remarkable progress was achieved in scaling simulations to unprecedented levels. The team successfully generated and utilized high-order meshes ranging from 130 million to 16 billion degrees of freedom for the Common Research Model High-Lift (CRM-HL) Case 2.4. The 760 million DoF simulations reached statistical convergence, producing accurate time-averaged lift and drag coefficients with high-resolution flow visualizations. Scalability tests were conducted using a 4.9 billion DoF mesh on the Mare Nostrum 5 supercomputer, with plans to scale to the full machine capacity of 4096 GPUs.

Advanced Mesh Generation and GPU Optimization

Significant advancements were made in high-order mesh generation for complex geometries, utilizing a combination of open-source tools (GMSH) and commercial software (ANSA, Pointwise). The team overcame substantial challenges in creating curved high-order meshes near highly curved wall boundaries. The implementation of OpenACC provided optimal GPU performance, enabling substantial energy consumption reductions while maintaining computational accuracy.

Combustion and Emissions Integration

A species model was successfully integrated into the SOD2D source code, establishing the framework for coupled Alya-SOD2D simulations to compute emissions formation. Preliminary simulations using representative aircraft geometries and relevant operating conditions were conducted, guided by thermodynamic principles from typical turbojet and turbofan engine equations.

Intelligent Post-Processing Infrastructure

The development of in-house POD and DMD tools enables sophisticated analysis of massive simulation datasets. The successful interfacing of Alya with the VISTLE visualization tool provides in-situ visualization capabilities, including Virtual Reality support, transforming how researchers analyse and understand complex flow phenomena in exascale simulations.

5.3 Impact and Significance

The achievements of EXCELLERAT P2 represent a paradigm shift in computational aerospace engineering, directly addressing critical industry needs for sustainable aviation. The project's impact extends across multiple dimensions of scientific and technological advancement.

Environmental and Regulatory Compliance

The development of integrated aircraft-engine emission modelling capabilities directly supports the aviation industry's transition toward meeting stringent EU decarbonisation targets. By enabling detailed predictions of both CO2 and non-CO2 emissions (including NOx and soot), the platform provides essential tools for designing next-generation low-emission aircraft systems.

Computational Efficiency and Energy Sustainability

The transition to high-order spectral element methods combined with GPU optimization represents a significant advancement in computational efficiency. The dramatic reduction in energy consumption while maintaining or improving accuracy aligns with broader sustainability goals in scientific computing. This efficiency gain is particularly crucial as simulations scale to exascale levels.

Industry Integration and Practical Applications

The project's participation in the High Lift Prediction Workshop (HLPW) demonstrates real-world applicability and industry relevance. The ability to simulate complex aircraft geometries with unprecedented fidelity provides immediate value to aerospace manufacturers seeking to optimize designs for aerodynamic efficiency and emissions reduction.

Scientific Methodology Advancement

The integration of artificial intelligence through deep reinforcement learning for flow control, combined with advanced data analytics using POD and DMD, establishes new methodological frameworks for understanding and controlling complex fluid phenomena. This approach transforms traditional CFD from a purely physics-based discipline to a hybrid physics-AI methodology.

HPC Architecture Utilization

The successful scaling to multi-billion degree-of-freedom simulations on modern GPU-accelerated supercomputers demonstrates effective utilization of cutting-edge HPC infrastructure, providing a template for other large-scale scientific computing applications.

5.4 Next Developments

The final phase of EXCELLERAT P2 will focus on consolidating achievements and pushing toward full system integration.

Simulation Scale-Up efforts will complete high-fidelity simulations using the 16 billion DoF mesh, with comprehensive performance evaluations scaling to the full Mare Nostrum 5 capacity. This will establish benchmarks for exascale CFD simulations in aerospace applications.

Coupled System Implementation represents a critical next step, involving rigorous testing of the integrated species model and deployment of the coupled Alya-SOD2D system for comprehensive emissions computation within complex flow fields. This integration will enable the first truly comprehensive aircraft-engine emission simulations at scale.

Advanced Flow Control Integration will expand beyond current capabilities to implement both classical and AI-based flow control approaches. The development of deep reinforcement learning algorithms for high-lift wing control will establish new paradigms for adaptive aerodynamic optimization.

Mesh Convergence Analysis will provide fundamental insights into numerical accuracy by analysing convergence behaviour as discretization order increases and conducting comprehensive mesh convergence studies for fixed-order, variable-element-count configurations.

Production Deployment activities will focus on transitioning research codes to production-ready tools suitable for industry adoption, including comprehensive documentation, user interfaces, and integration with existing aerospace design workflows.

The culmination of these developments will establish EXCELLERAT P2 as a foundational platform for next-generation sustainable aviation design, providing industry and researchers with unprecedented capabilities for understanding and optimizing complex aircraft-engine systems in the context of environmental sustainability and regulatory compliance.

6 UC-5 High-fidelity simulation of rotating parts

Partner: KTH, Code: Neko

6.1 Overview

The EXCELLERAT P2 project's Use Case 5 (UC-5) focuses on "High-fidelity simulations of rotating parts," specifically studying flow dynamics around drone rotors in hover configuration. This use case represents a continuation and enhancement of work initiated in EXCELLERAT Phase 1, addressing the rapidly expanding small Unmanned Aerial Vehicle (UAV) market, which is projected to reach USD 58.4 billion by 2026.

The research centers on investigating the "Iowa" rotor - a two-bladed, twisted rotor configuration - using the high-order Computational Fluid Dynamics (CFD) solver Neko with GPU acceleration. The primary objectives encompass performing high-fidelity simulations to analyse rotor thrust, wake characteristics, and noise generation. This work is particularly relevant given the increasing concern about noise pollution from frequent UAV operations in urban environments.

The project employs the Spectral Element Method (SEM) through Neko, a modern successor to Nek5000 that leverages object-oriented Fortran 2008 and provides superior GPU support through various hardware backends (CUDA, HIP, OpenCL). The comprehensive workflow spans the entire simulation pipeline from pre-processing through execution to post-processing, incorporating advanced techniques including Adaptive Mesh Refinement (AMR), aeroacoustic analysis, uncertainty quantification, and in-situ data processing.

6.2 Key Achievements

Advanced Meshing and AMR Implementation

The project successfully completed the development of high-order hex-based meshing capabilities, implementing conforming refinement schemes based on modified 3-refinement and 2-refinement approaches. A significant breakthrough was achieved in developing an improved blade model utilizing 2D-splines, replacing the previously insufficient chord scaling methodology. This advancement enabled the team to perform full AMR production runs using Nek5000, with results so compelling that visualization from a starting rotor simulation with Re=15000 was selected as a finalist for the ERCOFTAC Milton Van Dyke Flow Visualization Competition 2025.

Aeroacoustic Solver Development and Validation

The team successfully implemented and validated aeroacoustic capabilities by developing AcoNeko, a comprehensive Python package for sound computation, post-processing, and visualization. The validation process involved simulating sound radiation for a forward-facing step case at Re_H=8000, achieving unprecedented agreement with experimental data that had not been accomplished numerically before. This work was presented at the 51st Annual Meeting on Acoustics in Copenhagen, demonstrating the method's effectiveness for predicting noise from rotating blades.

Streaming Algorithms and Data Processing

The streaming algorithms task reached completion through collaboration with the Centre of Excellence in Exascale CFD (CEEC), resulting in the development of PySEMTools package. This achievement enables streaming/parallel Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD) analysis, addressing the critical challenge of performing these computationally intensive analyses on exascale simulations without requiring massive data storage.

Uncertainty Quantification Framework Extension

Significant progress was made in extending the uncertainty quantification framework to support the Neko solver through ADIOS2 interface integration. The team successfully implemented in-situ uncertainty quantification for first-order statistics, with the code-agnostic framework providing robust uncertainty estimation capabilities that can be applied across different simulation platforms.

In-Situ Visualization and Analysis Integration

The project achieved seamless integration of in-situ data analysis and visualization tasks through the PySEMTools package, enabling real-time extraction and analysis of simulation data from Neko using the ADIOS2 library. The implementation supports comprehensive visualization capabilities through standard Python packages including Matplotlib for 2D plotting and PyVista/VTK for advanced three-dimensional analysis tasks.

6.3 Impact and Significance

Scientific and Technical Impact

The achievements represent a significant advancement in high-fidelity simulation capabilities for rotating machinery, particularly in the context of UAV applications. The successful validation of aeroacoustic methods against experimental data fills a critical gap in computational aeroacoustics, providing researchers and engineers with validated tools for noise prediction in drone applications. The development of streaming algorithms addresses a fundamental scalability challenge in computational fluid dynamics, enabling analysis techniques previously impractical for large-scale simulations.

Industrial Relevance

The work directly addresses pressing industrial needs in the rapidly expanding UAV market, where noise pollution has become a significant concern for urban operations. The validated simulation capabilities provide manufacturers with tools to optimize rotor designs for reduced noise emissions while maintaining aerodynamic performance. The comprehensive workflow from meshing through analysis supports industrial design cycles by providing detailed insights into rotor performance characteristics.

Computational Science Advancement

The successful implementation of advanced features in Neko, including AMR and uncertainty quantification, demonstrates the potential for modern, GPU-accelerated CFD solvers to handle complex, industrially relevant problems. The integration of multiple analysis tools through standardized interfaces (ADIOS2) represents a significant step toward comprehensive, real-time simulation analysis capabilities that are essential for exascale computing applications.

Research Community Impact

The development of open-source tools like AcoNeko and PySEMTools contributes valuable resources to the broader computational fluid dynamics and aeroacoustics research communities. The validation against experimental data and participation in prestigious competitions demonstrates the scientific rigor and practical applicability of the developed methods.

6.4 Next Developments

Short-term Objectives (Year 3)

The immediate focus will be on finalizing the CPU version of the non-conforming Neko solver, which represents a critical milestone for the AMR implementation. The team plans to address spurious pressure fluctuations observed at non-conformal interfaces during rotating blade simulations with AMR, investigating various mitigation strategies to ensure accurate sound emission predictions.

Advanced Aeroacoustic Capabilities

Future development will focus on performing hybrid aeroacoustic simulations of the drone rotor, building upon the validated forward-facing step methodology. The team plans to couple Neko with external acoustic wave equation solvers to enable comprehensive near-field and far-field acoustic analysis. This will involve implementing sophisticated data interpolation schemes and validating results against available experimental data for the Iowa rotor configuration.

GPU Implementation and Performance Optimization

A major development goal is the implementation of the GPU version of the non-conforming Neko solver, scheduled for the final project year. This advancement will unlock the full potential of modern GPU architectures for large-scale AMR simulations, potentially enabling unprecedented simulation scales and resolution capabilities.

Enhanced Integration and Collaboration

The project will strengthen collaboration with Work Package 4 partners, particularly integrating Neko with the InSiDS tool developed by Fraunhofer partners. Additionally, plans include exploring enhanced visualization capabilities by investigating CATALYST2 integration for PySEMTools, potentially expanding real-time analysis and visualization options.

Extended Uncertainty Quantification

Development will continue on higher-order statistics uncertainty quantification, building upon the successful first-order implementation. This extension will provide more comprehensive uncertainty bounds for complex turbulent flow statistics, enhancing the reliability of simulation predictions for engineering applications.

The roadmap positions the project to deliver a comprehensive, validated simulation framework that addresses critical industrial needs while advancing the state-of-the-art in computational fluid dynamics and aeroacoustics.

7 UC-6 Active control for drag reduction of transonic airfoils

Partner: CINECA, Codes: STREAmS

7.1 Overview

The STREAmS Use Case (UC-6) within the EXCELLERAT P2 project represents a cutting-edge investigation into active flow control for drag reduction on transonic airfoils, addressing one of aviation's most pressing environmental challenges. With the International Civil Aviation Organization projecting that aviation's carbon footprint could triple by 2050, consuming a quarter of the world's remaining carbon budget, this research directly targets fuel efficiency improvements through advanced aerodynamic control.

The project focuses on streamwise-traveling waves (StTW) of spanwise wall velocity as an active flow control method. This technique involves prescribing wall velocity according to the formula w_w(x,t) = A sin(κ_x x - ωt), where the resulting Stokes layer interacts with near-wall turbulence structures to disrupt momentum streak regeneration cycles. Previous studies on flat-plate flows demonstrated up to 30% net skin-friction reduction, and recent DNS of supercritical airfoil sections showed not only decreased skin friction but also increased lift through downstream shock displacement.

The research employs a three-tiered investigation approach: comprehensive parameter exploration at low Reynolds number ($Re = 3\times10^5$), moderate-Reynolds evaluation at $Re = 6\times10^5$, and ambitious high-Reynolds demonstration at $Re = 1.2\times10^6$. This systematic approach aims to identify robust actuation settings that deliver net drag savings without compromising lift while elucidating the underlying physics across realistic flight conditions.

The STREAmS-2 solver represents a significant advancement in computational fluid dynamics capabilities, supporting Direct Numerical Simulation (DNS) on heterogeneous high-performance computing architectures. The solver's multi-GPU compatibility across NVIDIA, AMD, and Intel platforms ensures flexibility and adaptability to diverse European HPC ecosystems.

7.2 Key Achievements

Solver Development and Modernization

The team successfully refactored the previous FLEW solver into the modular, object-oriented STREAmS-2.1 framework, enabling seamless integration of multiple computational backends. In collaboration with Work Package 3, both CUDA (NVIDIA) and HIP (AMD) implementations were validated through preliminary DNS at $Re = 3\times10^5$. Additional backends for CPUs and Intel GPUs were also tested and validated. This multi-platform capability prevents vendor lock-in and enables adaptation to the most favourable market conditions and HPC system availability across Europe.

Comprehensive Parameter Study

A landmark achievement was the completion of an extensive parametric study comprising 28 DNS cases (2 uncontrolled and 26 controlled) at $Re = 3\times10^5$, $M\infty = 0.7$, and $\alpha = 4°$. This comprehensive sweep of control parameters - oscillation frequency ω, streamwise wavenumber κ_x, and forcing amplitude A - established a detailed map of the control parameter space. The results, published in a peer-reviewed journal, demonstrate both friction reduction and lift enhancement via downstream shock displacement, with extrapolated full-aircraft cruise-drag savings on the order of 9% for approximately 1% power penalty.

Workflow Automation and Integration

The development of a Python-driven workflow management system represents a significant step toward code usability and accessibility. This automated pipeline manages every stage from mesh generation through DNS execution, including RANS-derived turbulence metrics for grid optimization and SLURM job dependency management. The system reads airfoil coordinates, flow conditions, and control parameters, then automatically generates curvilinear meshes, performs preliminary RANS simulations, analyses flow statistics, and executes DNS with minimal manual intervention.

Technical Enhancements for High-Fidelity Simulation

Critical improvements were implemented to ensure simulation accuracy and stability. These include interpolation of converged RANS solutions onto DNS grids for physically consistent initialization, far-field relaxation zones to the RANS state, and embedded sponge layers to absorb acoustic disturbances and prevent spurious wave reflections. Grid generation was enhanced to eliminate metric discontinuities and realign the finest C-mesh regions with actual flow wakes.

Resource Acquisition and Infrastructure Development

The team was granted computing time of over 1.1 million node-hours through EuroHPC and other funding calls, including participation in the FASTER project with 1 million node hours on LUMI-G. This substantial computational allocation positions UC-6 to deliver its ambitious high-Reynolds objectives and execute the largest active-control DNS simulations to date.

7.3 Impact and Significance

Scientific Impact

The research represents a significant advancement in understanding active flow control mechanisms for realistic aeronautical applications. The comprehensive parameter study at Re = 3×10^5 provides the first systematic mapping of StTW control effectiveness on transonic airfoils, while the planned high-Reynolds simulations at Re = 1.2×10^6 will push the boundaries of DNS capabilities in transport-relevant regimes. The validation of drag reduction and lift enhancement mechanisms through downstream shock displacement offers new insights into shock-boundary-layer-control interactions.

Technological Advancement

STREAmS-2.1 represents a state-of-the-art DNS solver with unprecedented multi-platform GPU compatibility. The solver's ability to achieve performance levels that reasonably reproduce peak capabilities across diverse computing devices ensures broad accessibility and sustainability. The automated workflow system significantly reduces the expertise barrier for conducting large-scale parametric studies, potentially accelerating research adoption across the broader CFD community.

Environmental and Economic Implications

With cruise drag being nearly proportional to fuel consumption and skin-friction drag comprising approximately half of total aircraft drag, the demonstrated 9% potential cruise-drag savings could translate to substantial environmental and economic benefits. For the global aviation industry, such improvements could significantly contribute to carbon emission reduction goals while providing competitive advantages to aircraft manufacturers and operators.

European HPC Ecosystem Strengthening

The project's multi-platform approach and successful deployment across various EuroHPC systems demonstrates European computational independence and technological sovereignty. The solver's vendor-agnostic design ensures continued competitiveness regardless of market fluctuations and supports the strategic objectives of European HPC initiatives.

7.4 Next Developments

Immediate Research Priorities

The team will complete the moderate-Reynolds study by performing two to three additional StTW-controlled DNS at Re = 6×10^5, selected based on drag-reduction trends from the completed baseline runs and the 26-case campaign at Re = 3×10^5. These simulations will validate control benefit persistence under increased turbulence levels and provide crucial scaling insights.

High-Reynolds Frontier Simulations

The execution of initial high-Reynolds simulations at Re = 1.2×10^6 represents the project's most ambitious milestone. Minimum targets include one uncontrolled and one controlled DNS on meshes of approximately 2×10^{10} cells, with preliminary completion estimates in Q3 2025 for the uncontrolled case and Q4 2025 for the controlled case. These simulations will constitute the largest active-control DNS to date and critically test both physical performance and computational robustness.

Cross-Platform Validation and Optimization

Comprehensive cross-platform testing will compare solver throughput and numerical consistency across NVIDIA and AMD GPUs, with relevant CPU-only comparisons. This validation ensures predictable operation on target EuroHPC systems and guides final adjustments to MPI decomposition and I/O strategies for optimal performance scaling.

Knowledge Dissemination and Community Impact

Post-processing efforts will systematically extract mean, RMS, and spectral statistics to assemble aerodynamic coefficients into a comprehensive project database. A public release of the STREAmS-2 workflow repository is planned, complete with tutorials and sample scripts, available on GitHub. The accumulated findings will be disseminated through peer-reviewed publications, conference presentations, and community workshops.

Long-term Vision and Sustainability

The project establishes a foundation for continued research into active flow control at unprecedented scales, with potential extensions to full aircraft configurations and multi-disciplinary optimization frameworks. The robust, vendor-independent solver architecture ensures long-term sustainability and adaptability to evolving HPC landscapes, supporting continued European leadership in computational aerodynamics and environmental sustainability initiatives.

8 FA-1 Engineering design and digital twin of the first wall of a tokamak fusion reactor

Partner: UL, Codes: L2G, OpenFOAM, Raysect

8.1 *Overview*

The University of Ljubljana's FA-1 use case represents a ground-breaking effort to develop a digital twin of the first wall of tokamak fusion reactors, specifically targeting the WEST reactor at CEA Cadarache and the ITER experimental reactor. This ambitious project addresses one of the most critical challenges in nuclear fusion: real-time monitoring and control of reactor components during fusion operations.

The digital twin workflow integrates three sophisticated computational approaches executed in sequence: field-line tracing for magnetically conducted particles, thermal modelling using finite element and finite volume methods (OpenFOAM), and optical simulation through Monte Carlo ray-tracing algorithms (Raysect). This comprehensive approach enables accurate prediction of surface temperatures from infrared camera measurements, accounting for complex optical phenomena including reflections and stray light that can compromise diagnostic accuracy.

The use case directly supports fusion energy development by providing tools for optimal engineering design of reactor first walls and enabling real-time detection of overheating components during operation. This capability is essential for preventing damage to critical reactor components and ensuring safe, stable fusion reactions. The project builds upon established community codes while pushing them toward exascale computing capabilities, making it relevant for next-generation high-performance computing applications.

The workflow addresses a fundamental challenge in fusion diagnostics: determining actual surface temperatures from camera images when surfaces are highly reflective. Only complete digital twin simulations can properly account for reflected light from multiple surfaces, enabling accurate temperature determination crucial for reactor safety and performance optimization.

8.2 *Key Achievements*

Enhanced Physical Modelling Capabilities

Since the M12 milestone, the project has significantly advanced its physical modelling sophistication. The L2G field-line tracing code now incorporates magnetic ripple effects, simulating realistic magnetic field disruptions in the φ direction caused by factors such as magnetic coil misalignment and current variations. This advancement enables more accurate representation of heat flux patterns on reactor components, as demonstrated in comparisons with experimental WEST reactor data.

The thermal modelling component has been substantially enhanced to handle multi-material systems typical of fusion reactor components. ITER panels composed of tungsten (W), copper-chromium-zirconium (CuCrZr), and stainless steel can now be accurately modelled with their distinct thermal properties. The optical simulation capability has been upgraded to incorporate Stefan-Boltzmann law radiation modelling, providing more realistic surface emission calculations.

Large-Scale Case Development

Two comprehensive large-scale simulation cases have been developed, representing significant advances in complexity and realism. The WEST reactor case provides a complete experimental validation framework, enabling direct comparison with measured infrared camera data from the operational tokamak. The ITER case focuses on actively cooled first wall panels, particularly panel 4, which features 6,300 castellated cubes in a geometry comprising approximately 9 million triangles.

These cases demonstrate the workflow's capability to handle realistic reactor geometries with full material complexity, including detailed modelling of cooling systems essential for sustained fusion operations. The castellated panel design, with its intricate cube structures attached to coolant-carrying subcomponents, represents one of the most challenging thermal modelling scenarios in fusion engineering.

Computational Performance Breakthroughs

Remarkable scalability improvements have been achieved across all three simulation codes. The L2G field-line tracing code, originally limited to single-node execution, now supports distributed memory architectures through OpenMPI implementation. Performance results demonstrate exceptional strong scaling: execution time reduced from 16 minutes in serial to just 24 seconds using 5,000 cores.

Similarly, impressive results were achieved for OpenFOAM thermal simulations, with ITER panel 4 calculations reduced from 3,490 seconds in serial execution to 16 seconds using 1,200 processes. The Raysect optical simulation code showed substantial improvements, decreasing from 13,695 seconds to 1,350 seconds on 1,560 processes. These performance gains make previously intractable problems computationally feasible for routine use.

Workflow Integration and Automation

A comprehensive Python library has been developed to streamline the complex multi-code workflow, significantly reducing manual intervention and potential for user errors. This automation framework handles the intricate data format conversions between L2G output, OpenFOAM input, and Raysect processing requirements, implementing optimized data formats (HDF5, netCDF) for efficient inter-code communication.

The integrated workflow now supports seamless execution across high-performance computing environments, with automated job allocation and resource management. This development transforms what was previously a research-grade, manually intensive process into a production-ready simulation capability suitable for routine engineering applications.

8.3 Impact and Significance

Advancing Fusion Energy Development

This digital twin capability represents a crucial advancement toward practical fusion energy implementation. Real-time temperature monitoring of reactor first walls is essential for preventing component damage and ensuring sustained fusion reactions. The ability to accurately interpret infrared camera data, accounting for complex optical effects in the high-reflectance environment of fusion reactors, addresses a fundamental diagnostic challenge that has limited reactor operation and safety.

The validation against experimental WEST reactor data demonstrates the workflow's readiness for practical application, while the ITER case development ensures compatibility with next-generation fusion facilities. This dual validation approach provides confidence in the methodology's applicability across different reactor designs and operational conditions.

High-Performance Computing Innovation

The scalability achievements position this workflow as a flagship application for exascale computing. The demonstrated strong scaling across three distinct computational approaches (particle tracing, thermal finite volume methods, and Monte Carlo ray-tracing) showcases the potential for fusion simulation applications to effectively utilize next-generation supercomputing resources.

The performance improvements enable parameter studies and uncertainty quantification analyses that were previously impossible due to computational constraints. This capability is essential for optimizing reactor designs and operating procedures, potentially accelerating fusion energy development timelines.

Methodological Contributions

The integrated multi-physics approach demonstrates how community codes can be effectively coupled to address complex engineering challenges. The successful integration of OpenFOAM, L2G, and Raysect provides a template for similar multi-code workflows in other demanding applications requiring sequential execution of different physics models.

The development of automated workflow management tools contributes to the broader high-performance computing community by demonstrating effective approaches to complex simulation pipeline management. These contributions extend beyond fusion applications to any field requiring integrated multi-physics simulations.

8.4 Next Developments

Scalability Enhancement and Exascale Preparation

Future development will focus on further optimization for exascale computing environments. While current distributed memory implementations show excellent performance gains, additional work is needed to fully exploit next-generation computing architectures. Planned improvements include enhanced load balancing algorithms, optimized memory management strategies, and advanced parallel input/output operations.

Specific attention will be directed toward the optical ray-tracing component, which presents unique challenges for extreme-scale parallelization due to its Monte Carlo methodology. Development of hybrid parallelization strategies combining distributed memory and accelerator technologies will be explored to maximize computational efficiency.

Comprehensive Case Study Completion

The WEST reactor simulation will be expanded beyond the current bumper and divertor target modelling to encompass the complete reactor geometry. This comprehensive approach will enable full validation against experimental infrared imaging data and provide a complete digital twin capability for operational reactor monitoring.

The ITER case will be finalized with full integration of actively cooled first wall panel systems, ensuring accurate representation of coolant flow dynamics and thermal management systems. This completion will provide a validated predictive capability for ITER operations, supporting both design optimization and operational planning.

Synthetic Diagnostics and AI Integration

A major focus will be the development of synthetic diagnostic capabilities that accurately replicate infrared imaging system responses. This development will enable comprehensive validation studies and support the creation of AI-based inverse algorithms for real-time temperature determination from camera data.

The inverse problem solution - determining surface temperatures from camera images in real-time - represents the ultimate goal for reactor control applications. Integration of FPGA and accelerator technologies will be explored to achieve the computational speeds necessary for real-time reactor monitoring and control systems.

Industrial Application and Technology Transfer

Preparation for technology transfer to industrial fusion development programs will include documentation of best practices, development of user training materials, and establishment of support frameworks for broader community adoption. The workflow's maturation toward production-ready status will support its integration into industrial fusion reactor development programs, potentially accelerating the path to commercial fusion energy.

9 Conclusions

All use cases in EXCELLERAT P2 have made a significant leap forward in their high-performance computing capabilities within the aerospace and energy sectors. Through the execution of the six innovative use cases and one further application, this initiative has not only advanced computational methodologies but also addressed critical challenges related to sustainability and efficiency in aircraft design and hydrogen combustion technologies. Each use case has demonstrated substantial achievements, from enhancing scalability in computational fluid dynamics to developing sophisticated workflows that integrate artificial intelligence with traditional engineering practices. The successful implementation of automated mesh refinement, advanced optimization techniques, and multi-physics simulations underscores the project's commitment to driving technological advancements that align with European decarbonization goals.