# HORIZON-EUROHPC-JU-2021-COE-01

EXCELLERAT P2

**The European Centre of Excellence for Engineering Applications**

**Project Number: 101092621**

# D2.7
# Updated Report on the AVBP 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):** | Antoine Dauptain | CERFACS |
| | Mohamed Ghenai | CERFACS |
| **Approved by** | Executive Centre Management | 16.12.2025 |
| **Reviewer** | Tufan Arslan | BSC |
| **Reviewer** | Michael Wagner | DLR |
| **Dissemination Level** | Public | |

| Date | Author | Comments | Version | Status |
|---|---|---|---|---|
| 30-04-2025 | Mohamed Ghenai | Initial draft | V0.0 | Draft |
| 30-04-2025 | Antoine Dauptain | Internal review | V0.1 | Draft |
| | | Integration of BSC review feedback | V0.2 | Draft |
| | | Integration of DLR review feedback | V0.3 | Draft |
| | | Integration of DLR review feedback | V0.4 | Draft |
| | | Integration of BSC review feedback | V1.0 | Final |
| 04-12-2025 | Antoine Dauptain | Revised version for resubmission, Integration of EU review feedback | V2.0 | Final |

# List of abbreviations

| | |
|---|---|
| ASMR | Automated Static Mesh Refinement |
| BOI | Body Of Influence |
| CEA | Commissariat à l'énergie atomique et aux énergies alternatives |
| CFD | Computational Fluid Dynamics |
| CI/CD | Continuous Integration/Continuous Deployment |
| CoE | Centre of Excellence |
| CoEC | Centre of Excellence in Combustion |
| DGAC | Direction Générale de l’Aviation Civile |
| ECFD7 | 7th Extreme Conference on Computational Fluid Dynamics |
| GPU | Graphics Processing Unit |
| H2 | Hydrogen |
| HPC | High-Performance Computing |
| LES | Large Eddy Simulation |
| Q | Quarter |
| RMS | Root Mean Square |
| TUB | Technische Universität Berlin |
| UC | Use Case |
| WP | Work Package |

# Executive Summary

This document presents the progress made in the AVBP Application Use Case UC-2 within reporting period 1 and 2 covering the first 30 months of the EXCELLERAT P2 project. Based on the detailed roadmap of the workflow development defined in deliverable D2.1 “Use-Case Execution Roadmap” [4], the workflow of the use case is summarised and the achieved progress with respect to the defined workflow, objectives and success criteria are presented.

In summary, the workflow development for UC-2 has progressed according to the schedule defined in deliverable D2.1. Work has been performed on the individual tasks planned for the first two years of the project.

# Table of Contents

1 Introduction ........ 7
2 Objectives of the Use Case ........ 7
3 Workflow Description ........ 8
4 Progress achieved ........ 8
5 Conclusion ........ 15
6 References ........ 16

# Table of Figures

Figure 1: Automated Static Mesh refinement workflow example ... 8
Figure 2: UC-2 Workplan as of the first two years of EXCELLERAT P2 ... 8
Figure 3: Simplified generic geometry TRAPVORTEX. Top a): Dimensions of the geometry; Middle b): Crinkly cut showing the homogeneous coarse mesh, c) Bottom: solution on an adapted mesh 15x finer, showed by the combustion model “thickening” on the median cut, and gradient of the velocity on inner and outer cylindrical cuts. ... 10
Figure 4: Overall view of a dual swirl hydrogen injector (Top), and close up of the initial surface mesh of the swirler device (Bottom). ... 11
Figure 5: Final adapted mesh obtained using the workflow excellerat on a cold run (M4, bottom) compared to the reference mesh done independently from this study. (Mref, Top) ... 12
Figure 6 : Comparison of mean and RMS velocity profiles 2mm downstream of the swirler: PIV is the experiment, M4 is the final adapted mesh using ASMR, Mref is the reference mesh. ... 13

# 1 Introduction

Future gas turbines will switch from fossil fuels to hydrogen. This technology is a key component of the decarbonisation objectives of the EU which include the transition to alternative non-carbon fuel sources in the next 20 years.

Such a change requires a massive adaptation in the design of combustion chambers, because hydrogen is a unique fuel, burning much faster, diffusing much faster, exploding faster than all other conventional fuels. The schedule allowed for the design of these chambers is very tight, and companies will rely heavily on simulations to this end.

In this use case, we will streamline the workflows required for fully integrated Hydrogen combustion simulations using advanced High-Performance Computing (HPC) on leading-edge HPC architectures. Namely, a realistic industrial-grade Large Eddy Simulation (LES) of a hydrogen combustion chamber will be showcased on pre-exascale and exascale EuroHPC JU systems.

# 2 Objectives of the Use Case

This use case focuses on bringing the simulation workflow for Hydrogen (H2) combustion applications to exascale using the AVBP solver [1].

Multiple challenges need to be addressed to make this possible. First, as with many LES cases, the quality of the grid used for the computations is a critical component. With this in mind, the potential of automatic mesh refinement techniques has been explored during EXCELLERAT Phase 1. We developed the KalpaTARU (formerly known as TREEADAPT) [10] library for massively parallel mesh adaptation on tetrahedral grids. For this use case, KalpaTARU needs to be brought up to specifications on exascale-level systems. Secondly, the solver itself, AVBP, needs to support current and future exascale hardware namely accelerated architectures and arm-based systems given the current European roadmap. HPC readiness is mainly handled in Work Package (WP) 3 and the current status is available in D3.1 “Report on Exa-Enabling Methodologies” [2]. And third, the workflow automation system needs to be stressed and brought to the level of interacting with high-end HPC clusters. Automation of this workflow is handled with Lemmings [3], which was also developed during EXCELLERAT Phase 1 [4] and will be enhanced for exascale computing during Phase 2.

Of course, automated mesh generation is only possible with adequate mesh adaptation metrics, these have been developed with the Centre of Excellence (CoE) CoEC and are not the object of efforts in EXCELLERAT P2.

# 3 Workflow Description

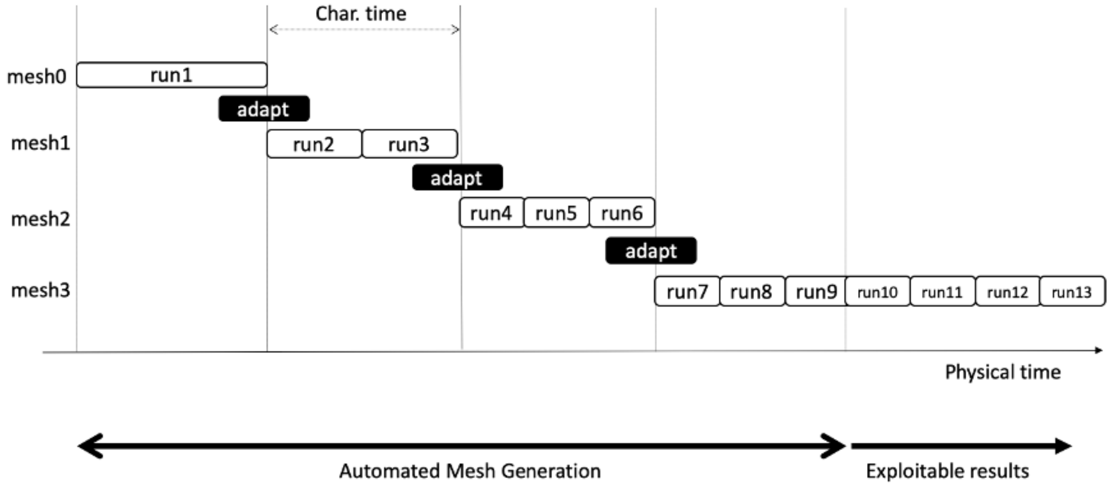

**Figure 1: Automated Static Mesh refinement workflow example**

In Figure 1, the Automated Static Mesh Refinement (ASMR) workflow under Lemmings is illustrated. The workflow proceeds as follows: we start from an initial mesh (mesh0) and perform a coarse LES simulation until convergence (run1). Using the data collected from this first simulation, we adapt the mesh to better capture the quantities of interest, creating mesh1 in the process. This updated mesh is then used to perform a new LES simulation, requiring run2 and run3 (a restart of run2 if necessary). Lemmings can automatically handle restarts, accounting for job duration limits on the target system.

This simulation-adaptation loop is repeated as many times as necessary. In short, Lemmings introduces a novel adaptive workflow capable of seamlessly transitioning between simulation and mesh adaptation steps. This adaptability is based on the statistical convergence of flow features and the simulation’s characteristic time.

The ASMR workflow schedules Computational Fluid Dynamics (CFD) runs of increasing resolution, followed by mesh adaptation runs based on the last solution. This iterative process allows the user to generate a high-resolution, CFD-driven mesh in a single operation.

For UC-2, the workflow has been customised for hydrogen combustion simulations using AVBP and mesh adaptation tools developed at CERFACS.

# 4 Progress achieved

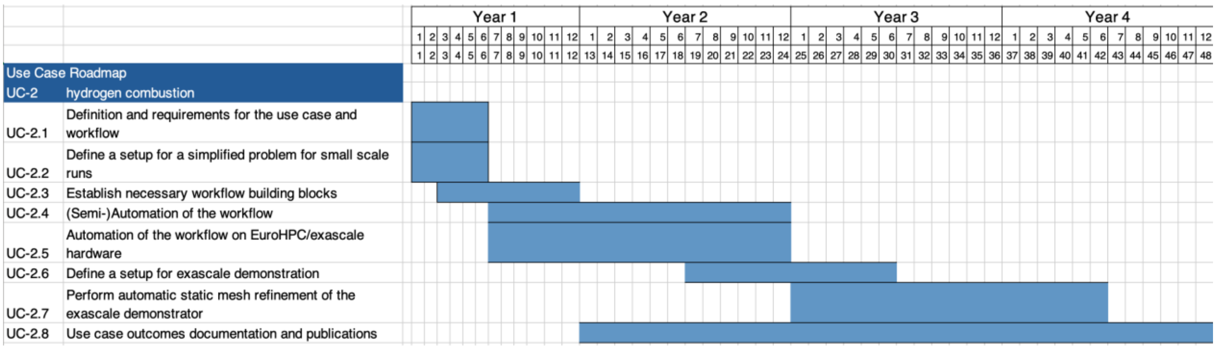

| | |
|---|---|
| Use Case Roadmap | |
| UC-2 | hydrogen combustion |
| UC-2.1 | Definition and requirements for the use case and workflow |
| UC-2.2 | Define a setup for a simplified problem for small scale runs |
| UC-2.3 | Establish necessary workflow building blocks |
| UC-2.4 | (Semi-)Automation of the workflow |
| UC-2.5 | Automation of the workflow on EuroHPC/exascale hardware |
| UC-2.6 | Define a setup for exascale demonstration |
| UC-2.7 | Perform automatic static mesh refinement of the exascale demonstrator |
| UC-2.8 | Use case outcomes documentation and publications |

**Figure 2: UC-2 Workplan as of the first two years of EXCELLERAT P2**

The first year of EXCELLERAT P2 concentrated on defining the requirements, the use case and the steps to follow to achieve our goal within the four years of the project.

In the second period of the project (13-30 months), this work has been extended, identifying new requirements and integrating them into the workflow construction. One such requirement was the systematic collection of job data points), which will be detailed in more detail in the second part of this section.

Figure 2 shows a detailed workplan with regard to Work Package 2, and the work is on schedule.

As mentioned above, it was decided to use the Lemmings [3] workflow environment as a follow-up to EXCELLERAT P1 and for P2, extend it to manage new cluster environments (such as EuroHPC JU systems).

So far, we have extended lemmings to the TOPAZE cluster from CCRT [11] for our tests in the first year. In Q4 2023, it was tested on LUMI C and in 2024 in LEONARDO at CINECA.

Concerning the target geometry, it was decided that we would focus on progressively complex cases using ASMR, rather than only the final H2 combustion industrial use case. This approach guarantees the generality and reusability of the workflow to other hydrogen combustion cases. Additionally, it eliminates any confidentiality concerns until after the Month 24 milestone when the exascale demonstrator tasks has required the real geometry from our industrial partner.

Lemmings on TOPAZE has therefore been used for the initial ASMR tests, with the TRAPVORTEX generic simplified geometry (Figure 3, left). It is a hollow channel of annular cross section, with a length of 0.2 m and an outer radius of 0.1m, for a total volume of 0.479 l. This geometry does not represent any real device. It was designed to run the lightest possible case while featuring all features found in actual aircraft combustion chambers. In particular, there is a pair of axi-periodic conditions, inlet slit only two-cells wide forcing the use of flux-based boundary conditions for cooling films, a wall for effusion cooling, and a jet in cross-flow. The discretisation is very coarse and homogeneous, with only 33 884 nodes and 177 563 tetrahedra of similar sizes. The mesh file is available in the open-source repository Gilgamesh [5].

Therefore, this case contains all the conditions of a real industrial Large Eddy Simulation within a small grid and the same modelling used for large industrial use cases, allowing for fast prototyping and validation (Q9). Figure 3 provide details of the case, as well as a view of a final simulation. Note how the jet-in-cross-flow wake is deviated clockwise due to a rotating motion of the main flow in the duct around its axis (Fig. 3c)

The initial tests to validate the mesh adaptation metrics (developed within the CoE CoEC and not included in this project) converged to a 2.5M element mesh with adequate flow feature representation. This “smoke test” was necessary before tacking real combustion devices.

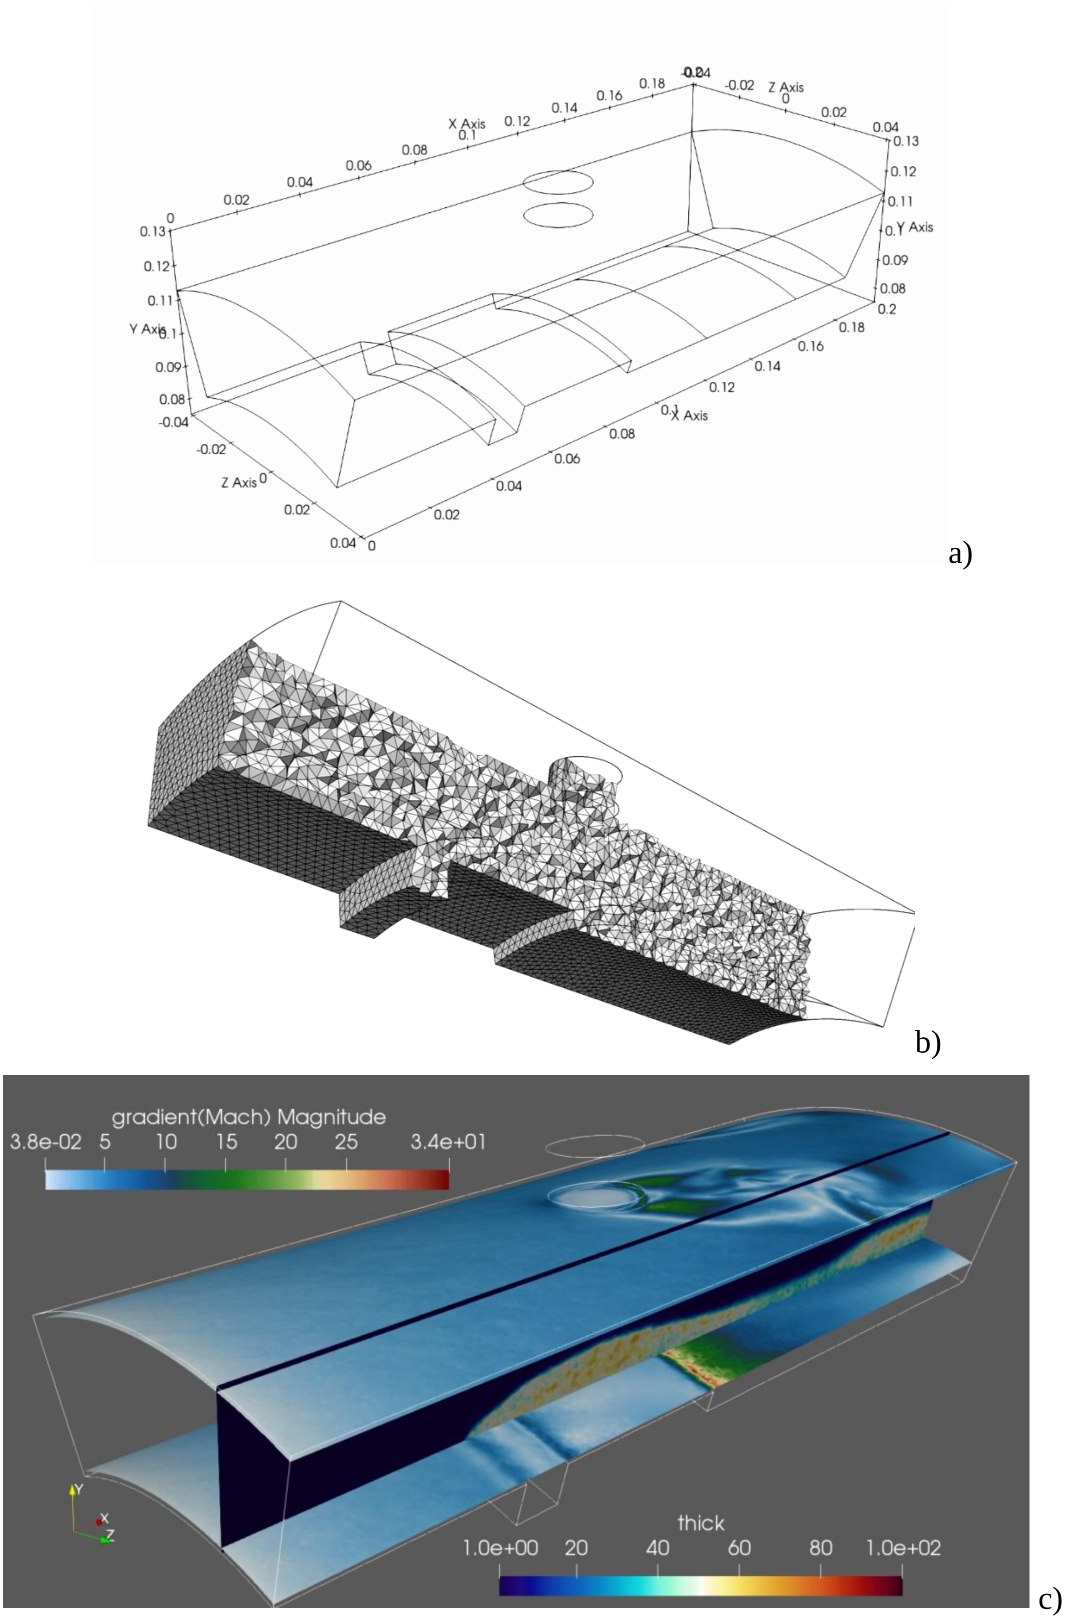

**Figure 3: Simplified generic geometry TRAPVORTEX. Top a): Dimensions of the geometry; Middle b): Crinkly cut showing the homogeneous coarse mesh, c) Bottom: solution on an adapted mesh 15x finer, showed by the combustion model "thickening" on the median cut, and gradient of the velocity on inner and outer cylindrical cuts.**

End-users from Safran Group tested the latest workflow and gave positive feedback since T12 of EXCELLERAT P2. A practical test of UC-2 was conducted during the one-week of the 7th Extreme Conference on Computational Fluid Dynamics (ECFD7) workshop (Q5). Several configurations were successfully run, including the AHEAD premixed H2/air injector [6], a laboratory-scale Hydrogen combustor developed at the Technische Universität Berlin (TUB).

Finally, engineers for Safran Aircraft tested the ASMR workflow on a more complex case, a dual swirl hydrogen injector experimentally studied in [7], [8].

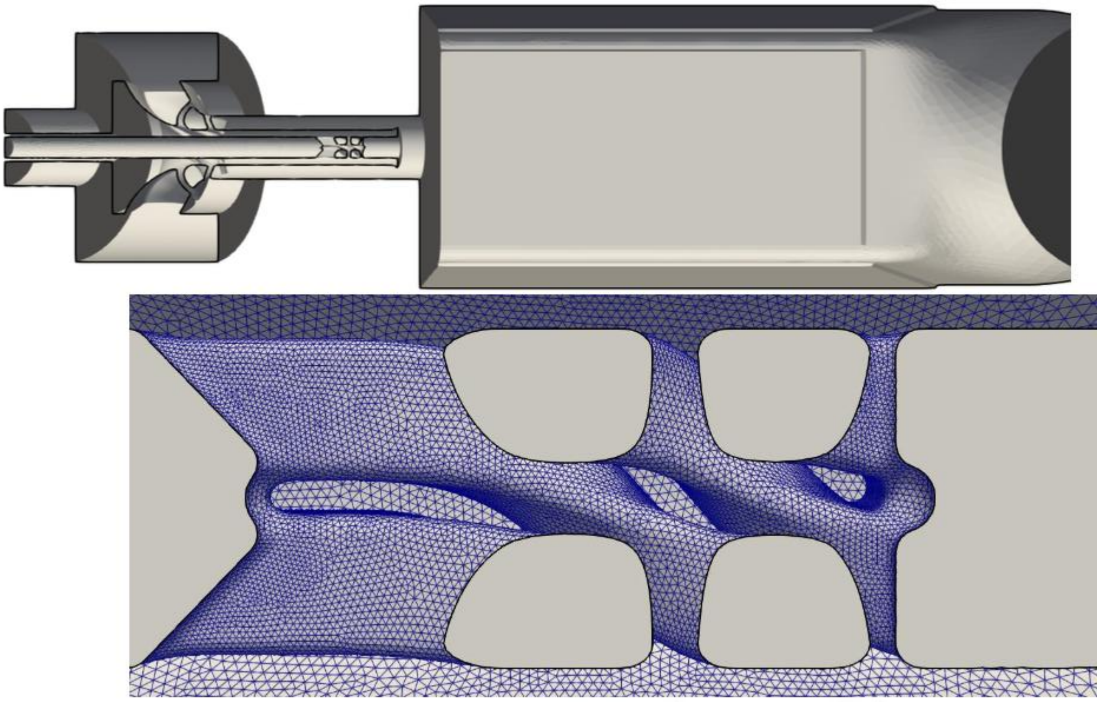

**Figure 4: Overall view of a dual swirl hydrogen injector (Top), and close up of the initial surface mesh of the swirler device (Bottom).**

This workflow application stands out because it was performed exclusively by actual end-users, Yahou Tarik and Nicola Detomaso from the company Safran Aircraft, within the project H2-tech funded by French Organism DGAC. The support from Cerfacs staff was negligible. The complexity of the configuration is clearly visible on the swirler shape, Fig. 4, which makes the manual positioning of Body Of Influence (BOI) very difficult.

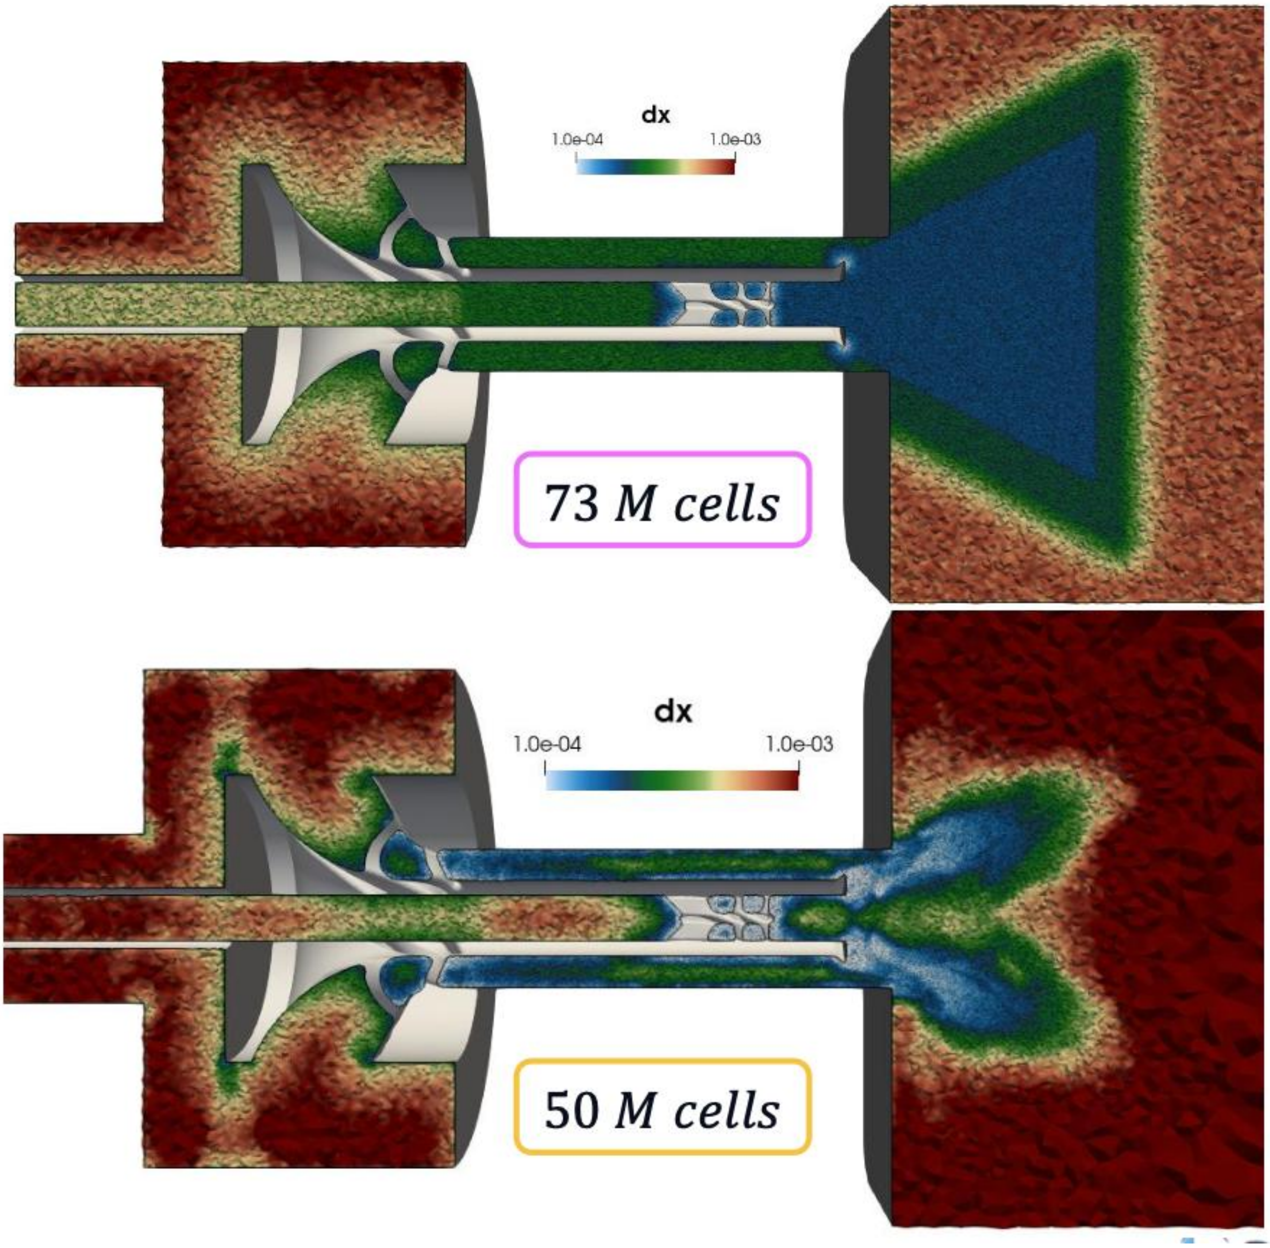

**Figure 5: Final adapted mesh obtained using the workflow excellerat on a cold run (M4, bottom) compared to the reference mesh done independently from this study. (Mref, Top)**

The workflow successfully produced good quality meshes with composite refinement on inner surfaces, driven by a Y+ sensor, and the shear zones of the swirl, driven by a sensor on Kinetic Energy Dissipation, after four sequences of run-and-remesh, hence the name “M4” (See Fig. 5).

Remarkably this M4 mesh features 50 Mcells compared to the 73 Mcells of the reference discretisation “Mref”. The refinement zone is smaller, and with higher resolution. As a result, a cold flow run takes 1814 CPUhours per millisecond of physical time on M4, while the same run consumes 3625 CPUhours / ms on Mref.

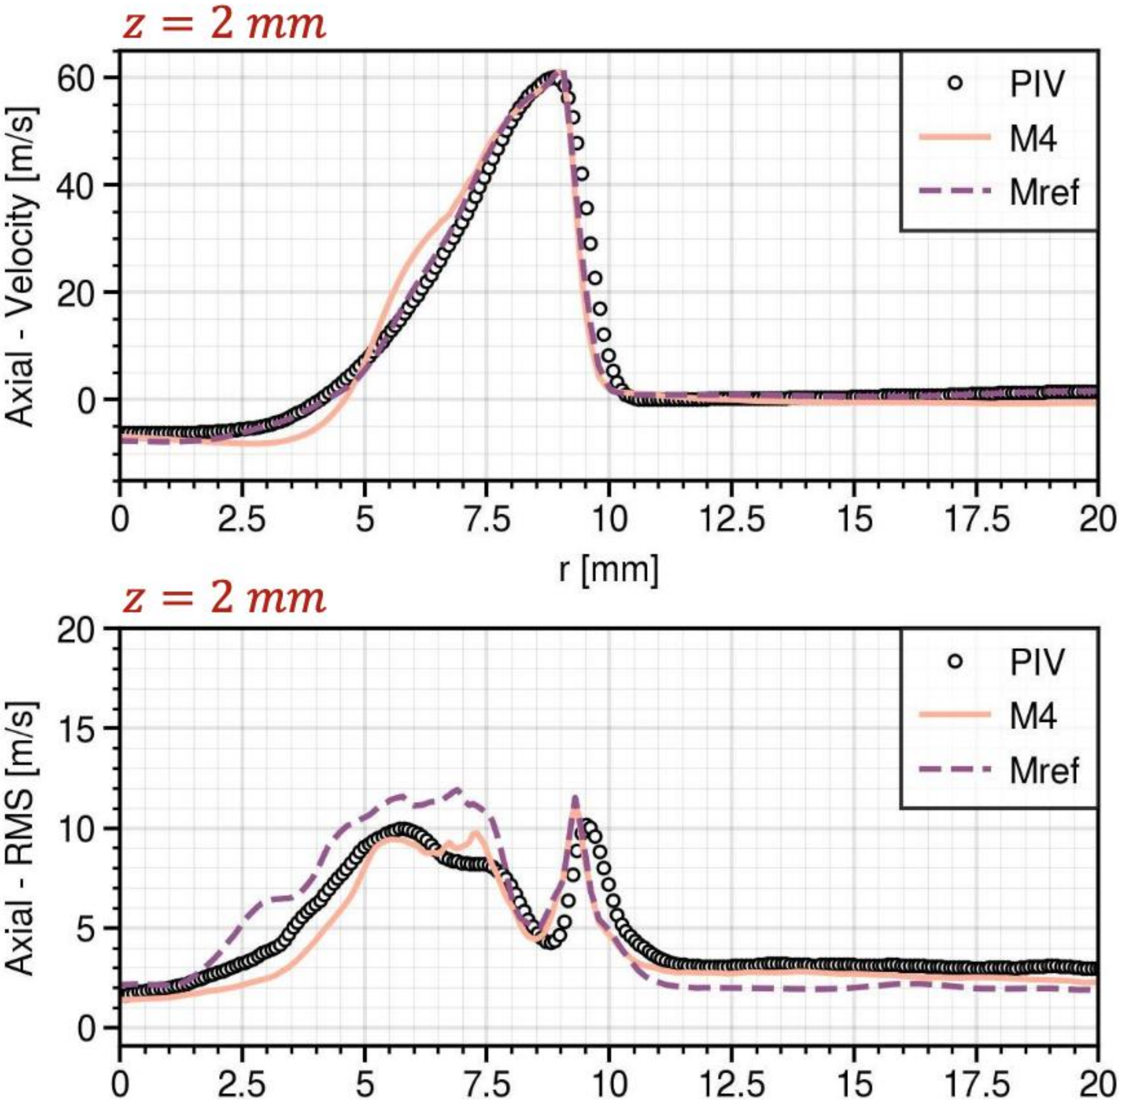

**Figure 6 : Comparison of mean and RMS velocity profiles 2mm downstream of the swirler: PIV is the experiment, M4 is the final adapted mesh using ASMR, Mref is the reference mesh.**

On the physical validation side, the flow features are correctly described with respect to the experiment. Both the reference flow and the one obtained by ASMR are reasonably matching the experimental data, as illustrated in Fig. 6. Both simulations are in good agreement, with a perfect match on the outer side of the swirl. The RMS values seem however closer to experimental data when using the adapted mesh M4.

As a result of these discussions, we implemented tools to collect job data points. AVBP now systematically records compilation options in its log output, improving future benchmark exploitability across different compilation setups. Initial tests of systematic job data points collection were conducted on both CERFACS's Kraken cluster and CEA's Topaze machine.

During this period, several improvements were made on the software stack backing the ASMR workflow:

Concerning KalpaTARU, the transition from the sequential mesh adaptation tool PyHIP to the massively parallel Kalpataru was also initiated during this period. The need for a "make periodic" tool was identified to address periodic mesh adaptation issues.

The ASMR workflow transitioned from the ParMETIS graph partitioner to PT-Scotch (Q6), resolving licensing concerns. A new Continuous Integration/Continuous Deployment (CI/CD) pipeline was established for testing the parallel periodic adaptation of the 2D LS89 cascade blade while continuing debugging and testing of the iterative parallel periodic AMR algorithm in 2D.

We can cite these improvements to Kalpataru (Q7):

- Resolved partitioning conflicts between AVBP and Kalpataru dependencies on ParMETIS/Scotch.
- Work on supporting periodic adaptation in Kalpataru continued.
- A prototype workflow integrating Kalpataru in a parallel job was developed.

The new workflow version that has been released contained (Q8):

- Adaptation task Tekigo using a generic script driven by a YAML configuration file, replacing ad-hoc scripts.
- Number and ratio of refinement steps computed automatically based on the initial and target meshes, rather than being user-defined.
- Prototype STL-to-initial mesh capability developed using GMSH.

# 5 Conclusion

Thus far, the ASMR workflow using Lemmings has been validated for meshes up to 50M elements. The state-of-the-art sequential mesh manipulation tool HIP sufficed for this stage. Meanwhile, WP3 advanced Kalpataru, focusing on reproducibility and scaling for exascale readiness.

Upcoming work will progressively replace HIP with Kalpataru when handling larger meshes (≥200M elements), a key step toward exascale compatibility.

Most validation cases to date have run on CPU platforms (Intel Cascade Lake and AMD Epyc Zen 2). WP3 has extended AVBP's GPU acceleration to models required for hydrogen combustion, and tests on GPU systems like LEONARDO followed.

The first two years of EXCELLERAT P2 demonstrated the feasibility and reliability of the ASMR workflow. Significant milestones have been reached, and the UC-2 workflow is nearing readiness for exascale demonstration.

Year 3 priorities will include:

- Scaling mesh adaptation to larger sizes using Kalpataru
- Strengthening GPU robustness and portability (with AMD GPU full support as a major milestone)
- Demonstrating this workflow on several EuroHPC systems.

# 6 References

[1] CERFACS, 2025. [Online]. Available: https://avbp-wiki.cerfacs.fr/.
[2] EXCELLERAT P2, Deliverable D3.1 "Report on Exa-Enabling Methodologies".
[3] CERFACS, 2025. [Online]. Available: https://lemmings.readthedocs.io.
[4] EXCELLERAT P2, Deliverable D2.1 "Use-Case Execution Roadmap".
[5] CERFACS, "Gilgamesh project on gitlab.com," 2018. [Online]. Available: https://gitlab.com/cerfacs/gilgamesh/-/tree/master/trapvtx.
[6] T. G. Reichel, S. Terhaar and O. Paschereit, "Increasing flashback resistance in lean premixed swirl-stabilized hydrogen combustion by axial air injection.," *Journal of Engineering for Gas Turbines and Power*, p. 137(7), 2015.
[7] H. Magnes, H. Marragou, A. Aniello, L. Selle, T. Poinsot and T. Schuller, "Impact of Preheating on Flame Stabilization and NOx Emissions From a Dual Swirl Hydrogen Injector," *J. Eng. Gas Turbines Power*, p. 051004, 2024.
[8] T. Capurso, D.: R. E. Laera and B. Cuenot, "NOx pathways in lean partially premixed swirling H2-air turbulent flame," *Combustion and Flame*, p. 112581, 2023.
[9] EXCELLEAT P2, 2025. [Online]. Available: https://www.excellerat.eu/.
[10] K. Mohanamuraly and G. Staffelbach, "Software, Kalpa-Taru," 2024. [Online]. Available: https://gitlab.com/cerfacs/kalpataru
[11] https://www-ccrt.cea.