HORIZON-EUROHPC-JU-2021-COE-01

EXCELLERAT P2

The European Centre of Excellence for Engineering Applications

Project Number: 101092621

D2.6 Report on the AVBP 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):	Gabriel Staffelbach		CERFACS
	Antoine Dauptain		CERFACS
Approved by	Executive Centre Management		13.11.2023
Reviewer	Herbert Owen		BSC
Reviewer	Michael Wagner		DLR
Dissemination Level	Public		

Date	Author	Comments	Version	Status
26-10-2023	Gabriel Staffelbach	Initial draft	V0.0	Draft
27-10-2023	Antoine Dauptain	Internal review	V0.1	Draft
30/10/2023	Gabriel Staffelbach	Integration of BSC review feedback	V0.2	Draft
7/11/2023	Gabriel Staffelbach	Integration of DLR review feedback	V0.3	Draft
8/11/2023	Gabriel Staffelbach	Integration of DLR review feedback	V0.4	Draft
8/11/2023	Gabriel Staffelbach	Integration of BSC review feedback	V1.0	Final

List of abbreviations

ASMR	Automated static mesh refinement
CFD	Computational Fluid Dynamics
COE	Center of Excellence
COEC	Center of Excellence in Combustion
LES	Large Eddy Simulation
WP	Work Package

Executive Summary

This document presents the progress made in the AVBP Application Use Case **UC-2** within reporting period 1 covering the first year of the EXCELLERAT P2 project. Based on the detailed roadmap of the workflow development defined in deliverable D2.1, 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 year of the project.

Table of Contents

Table of Figures

1 Introduction

Future gas turbines will switch from fossil fuels to hydrogen. This technology is a key component of the decarbonization 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 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 to exascale the simulation workflow for H2 combustion applications using the AVBP solver [1].

Multiple challenges need to be addressed to render this possible. First, as with many large eddy simulation (LES) cases, the quality of the grid used for the computations is a critical component. With this in mind, we explored during EXCELLERAT Phase 1 the potential of automatic mesh refinement techniques. We developed the KalpaTARU (formerly known as TREEADAPT) [4] 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 on D3.1. 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 [2], which was also developed during EXCELLERAT Phase 1[3] 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 CoE COEC and are not the object of efforts in EXCELLERAT.

3 Workflow Description

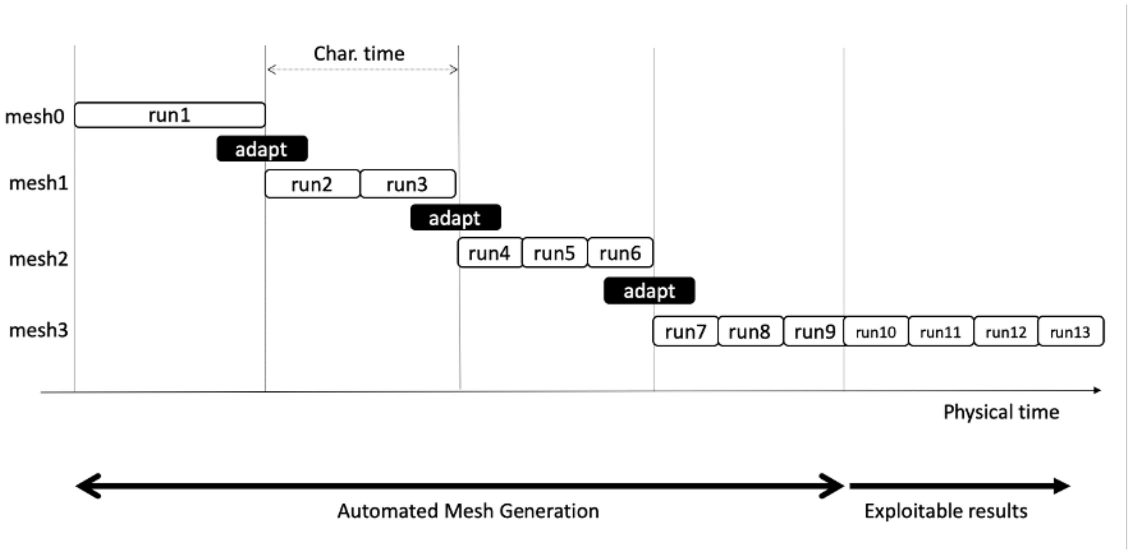

Figure 1: Automated Static Mesh refinement workflow example

In Figure 1, the ASMR workflow under Lemmings is illustrated.
The workflow is as follows: We start from an initial mesh (mesh0), and perform a coarse LES until convergence (run1). Then using the data collected during this first simulation we adapt the mesh to better represent the quantities of interest of the case creating mesh1 in the process. Mesh1 is then used to perform a follow-up LES until convergence, this requires run2 and run3 (a restart of run2). Lemmings is able to restart the simulation automatically taking into consideration the job duration limitations of the target system. Then we perform another adaptation. This loop is performed as many times as needed for the simulation. In short, Lemmings introduces a novel approach by enabling an adaptive workflow that seamlessly transitions between simulation and mesh adaptation steps. This adaptability is contingent upon the statistical convergence of flow features and the simulation's characteristic time. The ASMR workflow schedules CFD runs of increasing resolutions followed by mesh adaptation runs based on the last solution. This iterative process makes the user able to generate a high-resolution, CFD-driven large mesh in one single action.
For the current use case, this workflow is customised to the target hydrogen simulation using AVBP and the mesh adaptation tools from CERFACS.

4 Progress achieved

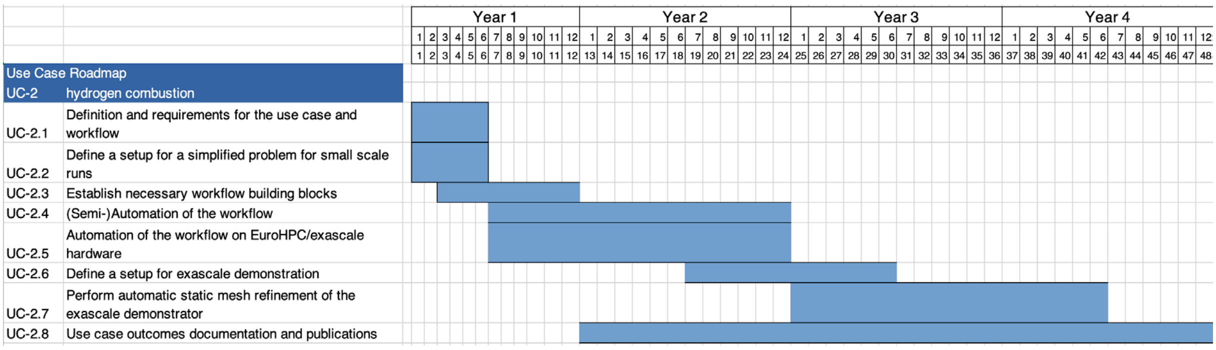

Use Case Roadmap	
UC-2	hydrogen combustion
UC-2.1	Definition and requirements for the use case and workflow
UC-2.2	Define a setup for a simplified problem for small scale runs
UC-2.3	Establish necessary workflow building blocks
UC-2.4	(Semi-)Automation of the workflow
UC-2.5	Automation of the workflow 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 1st year 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.
Figure 2 shows a detailed workplan with regard to Work Package 2 and we are on schedule.

As mentioned above, it was decided to use the Lemmings [2] 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 (https://www-ccrt.cea.fr/) for our tests within this period. In Q4 2023 it will be tested in 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 M24 milestone when the exascale demonstrator tasks will require the real geometry from our industrial partner.

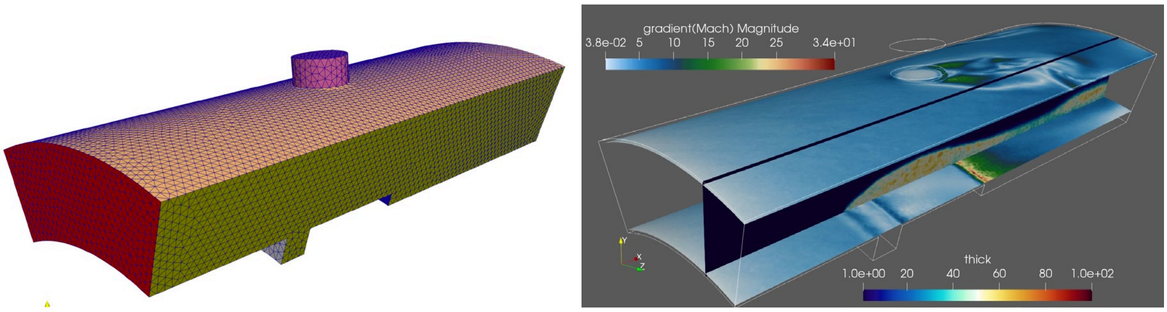

Figure 3: Left: Simplified generic geometry TRAPVORTEX; Right: Resulting instantaneous with of the ASMR workflow. Flame thickening and Mach gradient magnitude.

Lemmings on TOPAZE has therefore been used for the initial ASMR tests, with the TRAPVORTEX generic simplified geometry (Figure 3 left). This case contains all the conditions of a real industrial LES (complex boundary conditions, jet in cross-flow injection, annular shape) within a small grid (200.000 tetrahedra) and the same modelling used for large industrial use cases, allowing for fast prototyping and validation.
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 Figure 3 right.

Following this first validation, the workflow has been stressed on a more complex case, the TUB technically premixed H2/air injector [4], a laboratory-scale Hydrogen combustor developed at the Technische Universität Berlin. Figure 4 shows the consecutive mesh evolution during the ASMR workflow execution. From a basic 2.8M cell mesh (top image, coloured by edge size length), we progress to 5.7M, 11.5M, 23M and finally 46M elements. The progressive adaption is localised to resolve in detail the flame front and the main turbulent structures. It should be noted that the same simulation performed out of the box by an expert yielded an initial mesh of 200M elements and cost five times more.

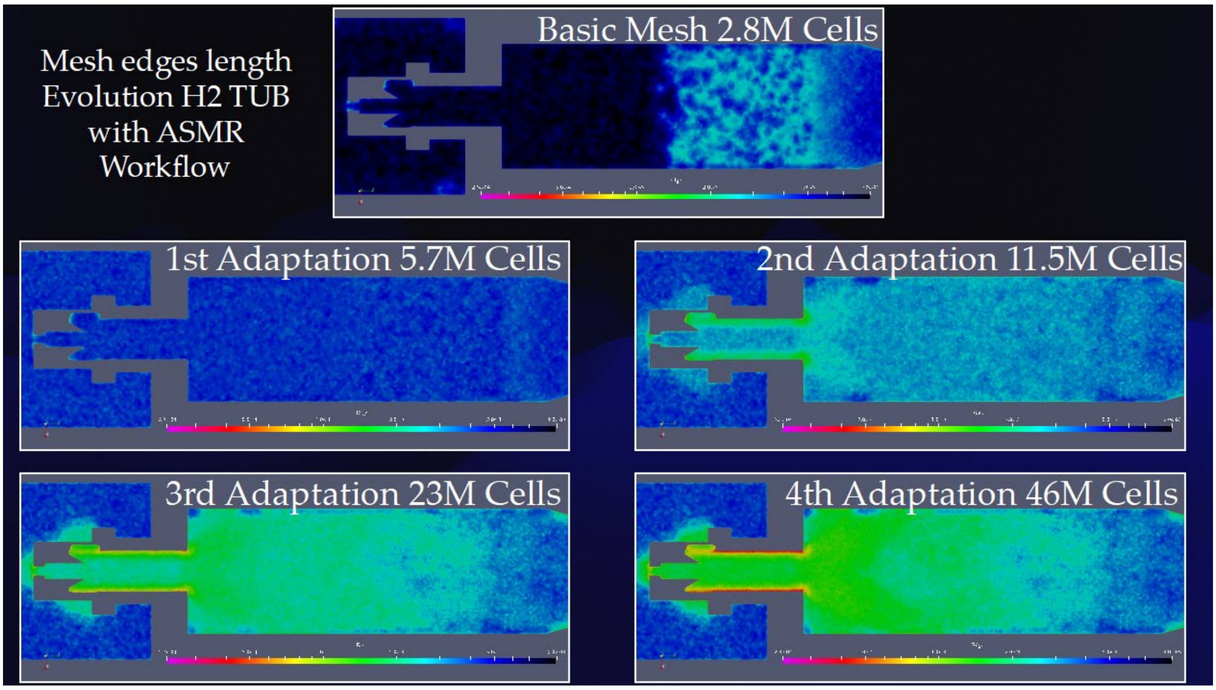

Figure 4: X=0 cut view of the TUB large eddy simulation mesh after successive mesh adaptation steps.

5 Conclusion

So far, these tests have focused on Lemmings's usage of the ASMR workflow. Since the size of the cases was rather small (maximum 50M elements) it was possible to use the current state-of-the-art sequential manipulation tool HIP [https://cerfacs.fr/coop/pyhip] instead of KalpaTARU. In the meantime, efforts on WP3 focused on the reproducibility of KalpaTARU. In the next phase, HIP will be replaced by KalpaTARU (when needed, i.e. for meshes larger than 200M elements) a requirement on the road to exascale compatibility.
These validation cases have been performed on CPU architectures (Intel Cascade Lake and AMD Epyc Zen 2). In WP3, GPU acceleration coverage of AVBP has been extended to the models required for H2 combustion and we will switch to GPU architectures such as LEONARDO for further testing.

The first year of EXCELLERAT has focused on the feasibility and reliability of the workflow. As evidenced in this document, multiple steps have been reached towards our goal and the ASMR workflow required for UC-2 is very near readiness for the exascale demonstrator. Year 2 will focus on enabling larger mesh adaptations with KalpaTARU and GPU robustness and portability (AMD GPU acceleration is a major step, see D3.1) as well as demonstrating this workflow on EuroHPC systems.

6 References

[1] AVBP code, https://www.cerfacs.fr/avbp7x
[2] LEMMINGS workflow Automator, https://lemmings.readthedocs.io
[3] EXCELLERAT Project, https://www.excellerat.eu/
[4] KalpaTARU library, https://gitlab.com/cerfacs/kalpataru

[5] Laera, E. Riber, B. Cuenot, “NOx pathways in lean partially premixed swirling H2-air turbulent flame”, In Proc. Combustion and Flame, Volume 248, 2023