

**HORIZON-EUROHPC-JU-2021-COE-01**




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

**Project Number: 101092621**

**D2.9**

**Updated Report on the m-AIA 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.

<b>Work Package:</b>	2	Use-Case Execution
<b>Author(s):</b>	Ansgar Niemöller, Matthias Meinke	RWTH
	Arno Feiden	FhG
<b>Approved by</b>	Executive Centre Management	3.06.2025
<b>Reviewer</b>	Dennis Grieger	USTUTT
<b>Reviewer</b>	Marvin Hubl	SSC
<b>Dissemination Level</b>	Public	

Date	Author	Comments	Version	Status
2025-04-29	Ansgar Niemöller, Arno Feiden	First draft based on deliverable D2.8 and progress since M12	V0.5	Draft
2025-05-16	Ansgar Niemöller	Second draft	V0.7	Draft
2025-05-26	Ansgar Niemöller	Final version	V1.0	Final

## List of abbreviations

ACARE	Advisory Council for Aeronautical Research in Europe
AI	Artificial Intelligence
APU	Accelerated Processing Unit
BO	Bayesian Optimisation
CAA	Computational Aeroacoustics
CFD	Computational Fluid Dynamics
CPU	Central Processing Unit
DG	Discontinuous Galerkin
DOF	Degrees of Freedom
FV	Finite Volume
FWH	Ffowcs Williams-Hawkings
GPR	Gaussian Process Regression
GPU	Graphics Processing Unit
HPC	High-Performance Computing
LES	Large-Eddy Simulation
M	Month
m-AIA	multi-physics AIA
MPI	Message Passing Interface
NACA	National Advisory Committee for Aeronautics
SMC	Small Metal Chevron
STL	Stereolithography
UC	Use Case

## Executive Summary

This document summarises the progress made since Month (M) 12 in the m-AIA application use case (UC) UC-3 within reporting period 2 covering the second year and half of the third year of the EXCELLERAT P2 project.

RWTH has advanced large-scale coupled Computational Fluid Dynamics (CFD) – Computational Aeroacoustic (CAA) simulations using the m-AIA framework, starting with a chevron nozzle simulation containing around 1 billion CFD cells and progressing to higher resolutions up to 3.7 billion cells to predict the jet flow and acoustic far-field noise spectra at Mach 0.9 and Reynolds number of 1 million. The corresponding CAA simulation involved up to 4.9 billion degrees of freedom. Far-field noise predictions were enhanced by improvements of the Ffowcs Williams-Hawkings (FWH) solver and it is found that the results are insensitive to the chosen parameters of the FWH surface. The largest simulation was executed on the Hawk High-Performance Computing (HPC) system using up to 4096 nodes achieving high efficiency, showing linear scalability compared to a baseline on 256 nodes. The results show that a full HPC system can be used effectively for aeroacoustic predictions with m-AIA, enabling the execution of large-scale use cases on a pre-exascale level. Workflow automation for m-AIA simulation pipelines was achieved through the development of a Python tool. Visualisation improvements were made using ParaView for handling large datasets and porting activities focused on adapting m-AIA for Graphics Processing Unit (GPU)/Accelerated Processing Unit (APU) architectures. An Artificial Intelligence (AI)-based optimisation workflow is in development in collaboration with project partner Fraunhofer SCAI (FhG). The modular approach supports parallel evaluation of candidate solutions, facilitating efficient optimisation processes for tasks like nozzle shape optimisation.

In summary, RWTH has made significant advancements regarding the various aspects relevant to the execution and objectives of UC-3.

## Table of Contents

1	Introduction .....	7
2	Objectives of the Use Case.....	8
3	Workflow Description.....	8
4	Progress achieved since M12 .....	10
5	Next Steps .....	16
6	References .....	17

## Table of Figures

Figure 1: Baseline (left) and typical chevron nozzle (right) .....	7
Figure 2: Computational setup with the LES or CFD, and the CAA domain for the jet noise prediction.....	7
Figure 3: Coupling of the CFD with the CAA solver and the FWH method.....	8
Figure 4: Workflow components for the constrained shape optimisation of chevron nozzles ..	9
Figure 6: Turbulent flow structures for jet flow from SMC001 chevron nozzle in chevron tip plane (top) and root plane (bottom) coloured by axial velocity .....	10
Figure 5: Aeroacoustic prediction for a jet emanating from SMC001 chevron nozzle. The visualisation shows the chevron nozzle and turbulent flow structures in the centre and the instantaneous acoustic pressure field in a plane. Colours indicate the magnitude of the flow velocity and grey scales the magnitude of the acoustic pressure.....	10
Figure 7: Comparison of DG and DG-FWH acoustic predictions at radial locations $r/D=8$ and $r/D=15$ .....	11
Figure 8: Comparison of acoustic far-field power spectral density at a distance of $100D$ , an angle of $0^\circ$ corresponds to the downstream direction. ....	11
Figure 9: Comparison of workload distribution of large-scale coupled CFD-CAA simulation on 256 and 4096 Hawk nodes .....	12
Figure 10: Scalability of large-scale CFD and coupled CFD-CAA simulations on Hawk.....	12
Figure 11: m-AIA DG benchmark with 1 billion DOF on Hunter: initial pctl-ported loops vs. rewritten compute kernels looping over individual nodes .....	13
Figure 12: Modular design of the optimisation framework .....	14
Figure 13: Best and worst airfoil shape found during one optimisation run and instantaneous flow field snapshots showing axial (top) and vertical (bottom) velocity components.....	15
Figure 14: Optimised scheduling where a second evaluation can be started without waiting for others to be finished .....	15

# 1 Introduction

In this use case, a shape optimisation of chevron nozzles is performed, which is representative for problems with expensive objective function evaluations in the important technical field of noise reduction. It includes a constraint of minimum jet thrust loss, which has to be weighted with the goal of noise reduction. The accurate prediction of the emitted sound is based on a turbulence scale resolving CFD method, directly coupled with a computational aeroacoustics solver, which is a typical example of a multiphysics application. With state-of-the-art solvers, the necessary large number of objective function evaluations can only be performed on exascale HPC systems to obtain results within acceptable time. An AI-based optimiser, developed by the project partner FhG, will be used to efficiently identify the optimal solution. The motivation for the use case is given e.g., by the Advisory Council for Aeronautical Research in Europe (ACARE), who established the Flightpath 2050, a new goal for more rigorous noise reduction by 65 percent relative to the capabilities of typical new aircraft in 2000. Despite the progressive introduction of high-bypass-ratio aircraft engines and chevron nozzles, which possess a sawtooth-like shape at the engine's trailing edge, jet noise still is a significant source of aircraft engine noise.

Unlike the optimisation of the aerodynamic performance or structural weight, noise reduction is still to a large extent an unsolved problem. One of the challenges connected to noise reduction is a reliable and accurate prediction of the sound pressure level in the far field, which is often generated by intricate flow phenomena. Therefore, turbulence scale resolving simulations have to be performed in many cases to obtain the correct sound pressure level. For example, the tip gap vortex in an axial fan can generate noise, which cannot be predicted with methods based on Reynolds averaged solutions. At present, shape optimisations are not possible for such cases due to the computational expensive evaluation of the objective function.

A typical chevron nozzle shape is shown in Figure 1 and the setup for the flow and the acoustic field prediction is depicted in Figure 2.




Figure 1: Baseline (left) and typical chevron nozzle (right)

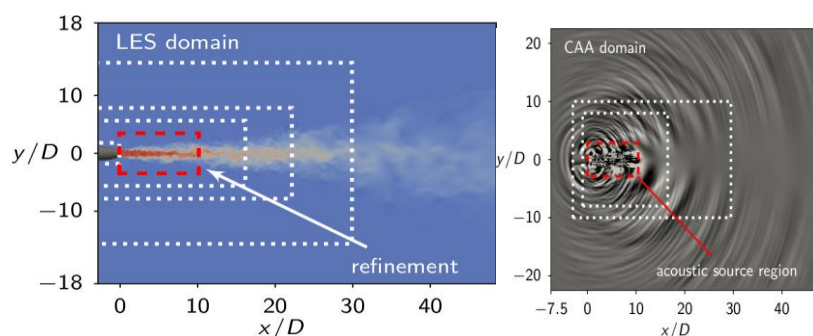


Figure 2: Computational setup with the LES or CFD, and the CAA domain for the jet noise prediction

## 2 Objectives of the Use Case


It is the objective of this use case to perform a shape optimisation of the chevrons at the nozzle exit to minimise the jet noise. The use of chevron nozzles, however, can lead to a substantial and unwanted loss of thrust during aircraft take-off and cruise flight. As changes in the chevron nozzle shape to minimise noise can lead to a loss of thrust, the optimisation has to be performed under the constraint that the thrust is only minimally reduced and its impairment remains below a given threshold. Since the objective function evaluation is computationally expensive, HPC systems of exascale have to be used efficiently. Therefore, an advanced workflow has to be designed, which can fully exploit the available computing resources. The goal is to implement all necessary workflow components such as the automatic generation of chevron shapes, the execution of the simulation for the flow field prediction and the determination of the complex objective function. Additionally, highly resolved simulations should be performed, which enable the identification of the essential noise source locations and to determine the effect of the various chevron parameters on the generation of the acoustic waves. For the execution of the use case, the workflow couples a high-fidelity 3-D aeroacoustics solver with state-of-the-art AI optimisation algorithms provided by the project partners FhG and data analytic tools to systematically perform thrust-constrained noise minimisation of chevron nozzle shapes.

For UC-3 the following success criteria are used during the course of the project to verify its success:

- Development of a workflow, which can automatically perform shape optimisations.
- Efficient usage of HPC hardware during the workflow execution.
- Identification of chevron shapes with minimised noise emission.
- Successful analysis of noise generation mechanisms based on large scale simulation runs on exascale hardware.
- Analysis of noise source mechanisms based on large scale simulations.

## 3 Workflow Description

The accurate prediction of the overall sound pressure level in the far field requires several numerical methods which must be combined to obtain the final result for the objective function. The CFD solver for the prediction of the turbulent flow field is directly coupled to the CAA solver for the determination of sound pressure levels in the acoustic near field, which then delivers unsteady data for a FWH method, which can compute the overall sound pressure level in the far field. The coupling between the CFD and CAA solver is schematically explained in Figure 3.




**Figure 3: Coupling of the CFD with the CAA solver and the FWH method**

A single simulation is performed in various stages to minimise the required computational effort. First, the Large-Eddy Simulation (LES)/CFD solver is executed alone, until a fully developed flow field and a sufficiently converged time averaged flow field is obtained. Subsequently, the CFD and CAA solver are executed in a fully coupled manner until the unsteady acoustic pressure signal is obtained in the near sound field. Finally, the history of the unsteady pressure signal on a closed control surface is used to predict the noise in the acoustic

far field. The overall sound pressure level is determined at reference locations for the evaluation of the objective function. In addition, the flow field is postprocessed to compute the nozzle thrust.

The full workflow to perform chevron shape optimisations is schematically depicted in Figure 4:



**Figure 4: Workflow components for the constrained shape optimisation of chevron nozzles**

In more detail, the workflow components for the jet noise prediction are:

- AI-based optimiser provides a chevron nozzle design, i.e., new chevron geometry parameters.
- Geometry generator generates Stereolithography (STL) surfaces for the mesh generation for the specific chevron nozzle shape.
- The grid generation generates refined meshes for the CFD and CAA simulation.
- The CFD simulation is performed to determine the time averaged mean flow field.
- The coupled CFD/CAA simulation is performed to predict the acoustic pressure in the near field.
- FWH method to predict the acoustic far-field.
- Postprocessing of the FWH and CFD results for the cost or objective function evaluation.


Regarding an optimisation process, the simultaneous execution of the many required individual runs poses a significant challenge for the coordination of the computational resources (various hardware platforms, queueing systems). Since the runs are expected to last for different lengths of time an intelligent, optimised coordination of the simulations is necessary. Since the AI-based optimisation algorithm is directly fed with data from the simulation (in situ) the simultaneous provision of large CPU/GPU resources is challenging. To enable a subsequent analysis of the simulation results, it is necessary to efficiently store the data on long-term storage systems which is challenging due to the large amount of data.

In summary, multiple simulations with different phases or configurations, various output files, and large data volume have to be performed. This requires a reliable and fault tolerant automation and high efficiency without the necessity of user interaction or supervision.


The exascale execution profile is defined by many small-scale runs for the shape optimisation, i.e., requiring on the order of  $O(1000)$  runs. In addition, a few large scale runs ( $O(10)$ ) will be performed for advanced data analytics to identify physical mechanisms. Pre-exascale hero runs ( $O(1)$ ) are planned for higher Reynolds number, coaxial jet configurations and very high-fidelity results to demonstrate exascale readiness.

## 4 Progress achieved since M12

Since M12 RWTH has made significant progress for UC-3 in the context of large-scale coupled CFD-CAA simulations using the multiphysics simulation framework m-AIA. Building upon an initial chevron nozzle simulation containing around 1 billion CFD cells a series of higher-resolution CFD simulations has been prepared. That is, three increasingly refined CFD grids are employed which contain around 1.5 billion, 2.1 billion and 3.7 billion grid cells to resolve the turbulent flow field of the jet flow from the Small Metal Chevron SMC001 chevron nozzle [1]. The flow conditions are a Mach number of 0.9 and a Reynolds number of 1 million based on the nozzle exit diameter. The corresponding coupled CFD-CAA simulation is performed with 4.9 billion degrees of freedom for the CAA simulation. These large-scale CFD and subsequent coupled CFD-CAA simulation runs have been executed to analyse the impact of varying CFD grid resolutions on the acoustic far-field noise spectra. To ensure comparability of acoustic power spectra across the different configurations the acoustic prediction for each case was run for a similar total simulation time.



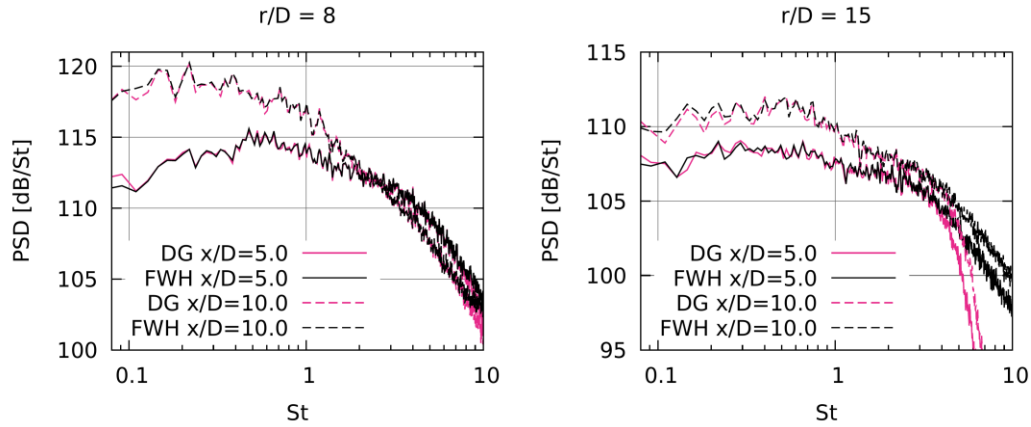
**Figure 6: Aeroacoustic prediction for a jet emanating from SMC001 chevron nozzle. The visualisation shows the chevron nozzle and turbulent flow structures in the centre and the instantaneous acoustic pressure field in a plane. Colours indicate the magnitude of the flow velocity and grey scales the magnitude of the acoustic pressure**



**Figure 5: Turbulent flow structures for jet flow from SMC001 chevron nozzle in chevron tip plane (top) and root plane (bottom) coloured by axial velocity**

A visualisation from a high-resolution simulation is given in Figure 5 showing the nozzle and turbulent flow structures and the instantaneous acoustic pressure field. A closer view of the jet flow development and the small-scale turbulent structures is given in Figure 6. In order to enable such visualisations, efforts have been dedicated to improving large-scale data visualisation and optimizing performance using ParaView for m-AIA solution files with more than 2 billion cells/data points.


Far-field noise predictions were obtained for all cases using the FWH solver using CAA data sampled on a high-resolution surface in the acoustic near-field. An important development was the identification and correction of an error in the formulation of the implemented equations within the FWH-solver, which had previously led to erroneous results depending on the specific case. To further improve the noise prediction workflow, the FWH solver was extended by



**Figure 7: Comparison of DG and DG-FWH acoustic predictions at radial locations  $r/D=8$  and  $r/D=15$**

implementing and testing Welch’s method for the power spectral density estimation of the acoustic far-field signal. In Figure 7 the Discontinuous Galerkin (DG) and DG-FWH predictions in the near far-field are compared for each two axial locations. As evident the solutions match nearly perfectly for the closer radial distance of  $r/D=8$  to the jet axis, up to which the DG solution representation still offers a sufficiently high spatial resolution to resolve nearly the full shown frequency range. By increasing the radial distance to the jet axis, the frequency cut-off due to coarser grid levels becomes evident in the DG solution.

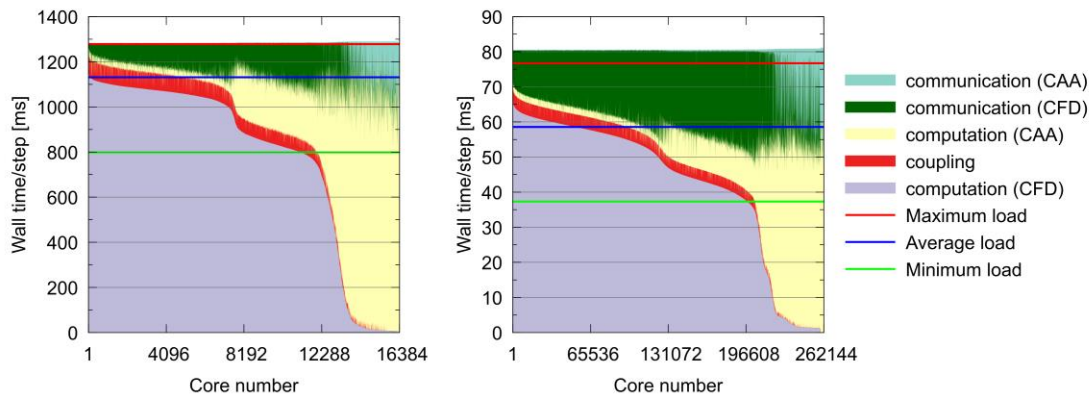
In Figure 8 the predicted power spectral densities in the far-field at a distance of  $100D$  are compared for two observer angles against reference data derived from experimental consensus as well as numerical results from Stich et al. [2]. The results of the high-resolution simulations demonstrate a strong agreement with the reference data across a wide frequency range. Overall, minimal differences have been observed between the far-field noise predictions obtained using the two highest CFD grid resolutions. Furthermore, the analysis of varying the FWH surface location and its spatial resolution showed that far-field noise predictions are insensitive to the chosen parameters, i.e., confirming a sufficiently high spatial resolution and independence of the approach on a meticulous surface placement.



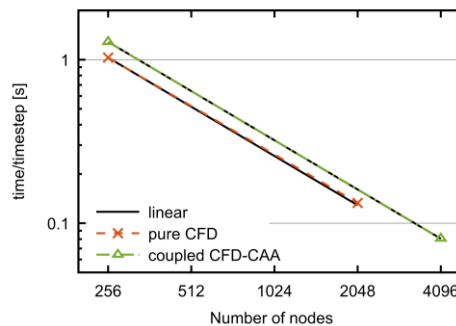
**Figure 8: Comparison of acoustic far-field power spectral density at a distance of  $100D$ , an angle of  $0^\circ$  corresponds to the downstream direction.**

## Large-scale simulations on pre-exascale level

In collaboration with HLRS the largest-scale CFD simulation of 3.7 billion cells was initially tested and then run on 2048 nodes of the HAWK system. Furthermore, the corresponding coupled CFD-CAA simulation was run in an XXL session on the full HAWK system using 4096 nodes totalling to 524k CPUs. The workload distribution for the coupled case is compared in Figure 9 on 256 and 4096 nodes, using 64 MPI processes per node. The dynamic load balancing approach in m-AIA is used to distribute the workload among processes. At high node counts it becomes increasingly difficult to eliminate all imbalances, due to very small partition sizes, while simultaneously the per-process performance benefits from a reduced memory footprint resulting in improved cache usage. Figure 10 shows the scalability of the large-scale simulations, compared to baseline runs on 256 nodes a nearly linear speedup is obtained on 2048 and 4096 nodes of Hawk, i.e., a full HPC system can be used effectively for aeroacoustic predictions with m-AIA. The results show that the numerical approach is highly scalable and it enables the execution of large-scale use cases on a pre-exascale level.



**Figure 9: Comparison of workload distribution of large-scale coupled CFD-CAA simulation on 256 and 4096 Hawk nodes**



**Figure 10: Scalability of large-scale CFD and coupled CFD-CAA simulations on Hawk**

## Workflow automation

As outlined in Section 3 the workflow for a single function evaluation and jet noise prediction, consists of several stages, each with multiple steps, which can be categorised in preprocessing, simulations, and postprocessing. which involve different tools, configurations, and requirements. As such, a basic automation of the m-AIA simulation pipeline has been developed to ensure ease of use, adaptability, and extensibility. The automation approach can be summarised as following. The user creates a baseline/template version of a m-AIA simulation setup including *toml* configuration files and jobscripts. Additionally, the workflow definition is captured in *json* format. The developed small tool written in Python automates the

simulation pipeline by reading the workflow definition and a *json* status file, checking the current step or running job, making necessary configuration and jobscript changes, submitting new jobs, and updating the status file accordingly. In summary, this makes the integration of all workflow steps and components into an optimisation loop feasible.

### Porting Activities and Code Optimisation

A key focus initiated during the second project year involved porting activities aimed at adapting the multiphysics simulation framework m-AIA to GPU/APU architectures to facilitate future use case executions and runs of the envisioned optimisation workflow on different types of HPC hardware. Apart from transforming code loops to C++ parallel-stl versions, the current improvements including code restructuring, clean-up and memory reduction are expected to enhance the execution efficiency of m-AIA not only on GPU/APU-based systems but also on traditional CPU-based HPC platforms. For the present use case UC-3 the Finite Volume (FV)-CFD and the DG-CAA solver of m-AIA and their respective coupling need to be ported. Since the DG method is based on a polynomial solution representation in each cell or element of the grid, with a potentially high number of Degrees of Freedom (DOF) inside a single element, compute kernels need to be rewritten to loop over individual nodes instead of full elements (cf. [3]). Through this approach the high computational effort required for each element can be split into small parts to be distributed among many threads. Figure 11 compares the performance in achievable time steps per second of the initially ported parallel-stl loops iterating over elements and surfaces to a first version of the rewritten kernels looping over individual nodes. The benchmark case consists of about 1 billion DOF and is executed on the AMD MI300A based Hunter HPC system at HLRS with each node containing 4 APUs. As evident, the changes to the loop structures result in a speedup of factor 6, while for 64 nodes/256 APUs still a parallel efficiency of about 70% is achieved.

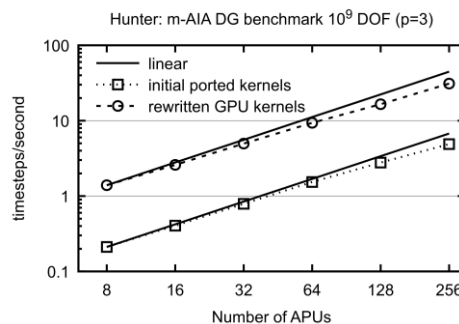


Figure 11: m-AIA DG benchmark with 1 billion DOF on Hunter: initial pstl-ported loops vs. rewritten compute kernels looping over individual nodes

### Optimisation workflow

Regarding the targeted AI-based optimisation, the collaboration between RWTH and FhG on the optimisation approach continued with recurring meetings and regular exchanges. Based on the simplified small-scale optimisation test setup consisting of a two-dimensional unsteady flow around a NACA profile at low Reynolds number, an optimisation workflow at low computational cost for the objective function evaluation was set up. The computational grids consist of approximately  $10^5$  cells, which combined with a quick convergence of the simulation renders the setup suitable for testing the optimisation workflow.

This CFD-based simulation model was integrated into a surrogate-based sequential closed-loop optimisation procedure, that is fully automated, parallelised, and can be restarted, see Figure 12. One optimisation strategy implemented is Bayesian optimisation (BO) with Gaussian Process Regression (GPR) [4] prior as the surrogate machine learning model. Due to the

modular approach of the optimisation procedure, another optimisation strategy based on a similar approach, implemented in a proprietary software developed by Fraunhofer SCAI, is also currently being tested. This software also supports the inclusion of simulation data in the optimisation process. Not just the relationship between design variables and optimisation target, but also the intermediate data of the simulation itself is then taken into consideration. This method is in development and tested on the NACA airfoil testcase and will later be applied on the chevron nozzle Results of an exploratory optimisation run of the NACA airfoil testcase are given in Figure 13, which shows the best and worst performing airfoil shapes discovered during

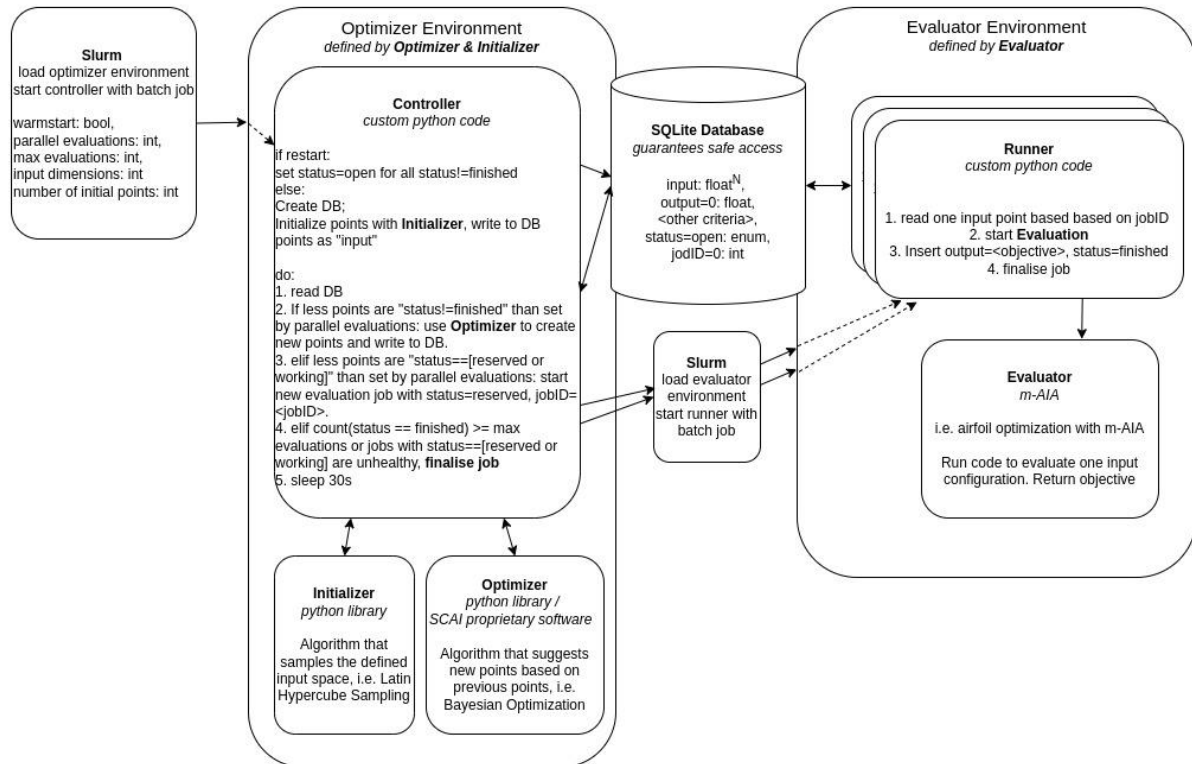
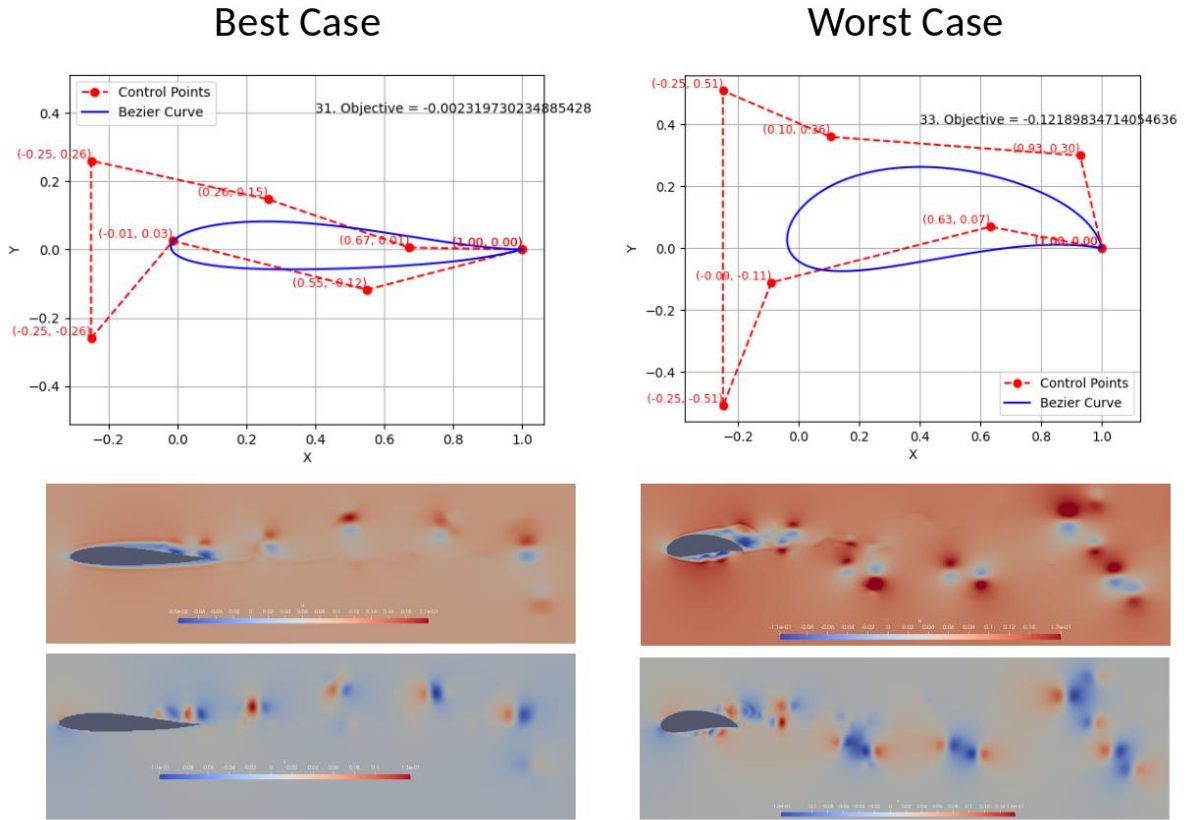


Figure 12: Modular design of the optimisation framework

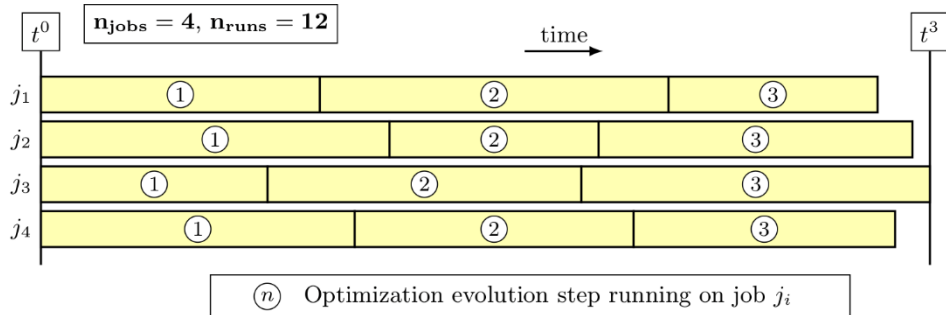
the optimisation process. The optimisation problem in this case is still considered a black box which expects as input seven 2-dimensional points forming a Bezier curve that defines the airfoil shape. Due to geometric constraints, this is reduced to a 9-dimensional input  $x$ . The output and optimisation variable  $f(x) = y$  is defined as a weighted sum of drag, lift, and thickness of the airfoil leading to a single-objective optimisation problem, i.e.,  $f(x) = c_D - 0.1 * c_L - 0.05 * A$ , with  $c_D$  and  $c_L$  the drag and lift coefficient, and  $A$  the airfoil cross-section area. As a starting point, the optimisation selected 20 points using Latin hypercube sampling, and then sequentially a new point is chosen based on the experience of every previous point with a Bayesian optimiser.

This entails building a surrogate model  $\hat{f} \approx f$ , a Gaussian process. A trade-off between exploration and exploitation is implemented by optimising an acquisition function  $\alpha$  instead of directly optimising the surrogate model  $\hat{f}$ , balancing expected improvement of the objective and the certainty of the model. In anticipation of the actual use case, the optimisation has been run for an additional 80 points, resulting in a total of 100 function evaluations.



**Figure 13: Best and worst airfoil shape found during one optimisation run and instantaneous flow field snapshots showing axial (top) and vertical (bottom) velocity components**

For the optimiser, several specific features are work in progress. This includes handling known constraints, like box constraints  $x_{min} \leq x \leq x_{max}$ , resp. the feasible set  $\hat{X}$  in general and unknown constraints, like certain configurations resulting in geometries that are not feasible. A fully modular design further enables the use of different environments for the optimiser and the simulation. This interplay between optimisation and evaluation has been set up to support the “embarrassingly parallel” evaluation of several candidate solutions at the same time. The pattern of candidate solutions that are pending enables even to suggest new candidates while evaluation of other candidates is still running, see Figure 14. With this, the seamless parallelisation is fully unlocked.



**Figure 14: Optimised scheduling where a second evaluation can be started without waiting for others to be finished**

## 5 Next Steps

Based on the conducted simulations for the SMC001 chevron nozzle, RWTH already commenced work on large-scale aeroacoustic simulations for the SMC000 round baseline nozzle to compare its flow and acoustic field data in a comparative study. Conducting comprehensive data analysis might provide valuable insights for future design optimisations and the noise source mechanisms. Simultaneously, efforts will focus on continuous improvements of the workflow components and the m-AIA simulation framework, with primary focus to complete the GPU/APU porting of the coupled FV-DG solvers. That is, based on a first version supporting more complex simulations on the MI300A based Hunter HPC system at HLRS, the performance will be further optimised. These developments will lay the foundation for the exascale readiness of the m-AIA code across compute architectures and enable the execution of use case shape optimisation runs with high efficiency in the near future. Preparations to perform an initial sample of smaller scale chevron nozzle simulation runs will continue, which will constitute the basis for optimisation attempts using the optimiser from FhG. The development of the optimiser continues to support satisfying requirements emerging from the use-case such as known and unknown constraints that are uncovered in the close corporation of RWTH and FhG. To improve performance in the final run, FhG primarily looks at grey-boxing as flashed out in chapter 4 and selecting established methods that have a record of good performance in similar problem setups, for example as analysed in the black-box optimisation benchmark [5].

## 6 References

- [1] J. Bridges and C. Brown, “Parametric testing of chevrons on single flow hot jets,” AIAA Paper 2004–2824, 2004. doi: 10.2514/6.2004-2824.
- [2] G.-D. Stich, et al. Jet Noise Prediction for Chevron Nozzle Flows with Wall-Modeled Large-Eddy Simulation. AIAA SciTech Forum 2021. doi: 10.2514/6.2021-1185.
- [3] M. Kurz, D. Kempf, M. Blind, P. Kopper, P. Offenhaeuser, A. Schwarz, S. Starr, J. Keim, A. Beck. (2025). GALÆXI: Solving complex compressible flows with high-order discontinuous Galerkin methods on accelerator-based systems. *Computer Physics Communications*. 306. 109388. doi: 10.1016/j.cpc.2024.109388.
- [4] Rasmussen, C. E.: *Gaussian Processes for Machine Learning*. In: Rasmussen, C. E., Williams, C. K. I., The MIT Press, Cambridge, Massachusetts, USA, ISBN 0-262-18253-X, 2006.
- [5] Santoni, M. L., Raponi, E., De Leone, R., and Doerr, C., “Comparison of High-Dimensional Bayesian Optimization Algorithms on BBOB”, arXiv:2303.00890, 2023.