Ref. Ares(2021)5370636 - 31/08/2021

ACROSS

HPC Big DAta ArtifiCial intelligence cross Stack PlatfoRm TOwardS ExaScale

D2.1 – Summary of pilots co-design requirements

Deliverable ID	D2.1
Deliverable Title	Summary of Pilots co-design requirements
Work Package	WP2
Dissemination Level	PUBLIC
Version	3.0
Date	2021 - 08 - 31
Status	Submitted
Deliverable Leader	CINECA
Main Contributors	AVIO, MORFO, ECMWF, SINTEF, IT4I, LINKS

Disclaimer: All information provided reflects the status of the ACROSS project at the time of writing and may be subject to change. This document reflects only the ACROSS partners' view and the European Commission is not responsible for any use that may be made of the information it contains.

The ACROSS project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955648. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Italy, France, Czech Republic, United Kingdom, Greece, Netherlands, Germany, Norway.

Published by the ACROSS Consortium

Document History

Version	Date	Author(s)	Description
0.1	2021-01-15	LINKS	First version
1.0		CINECA, AVIO, ECMWF, SINTEF, IT4I, LINKS	Working version
1.1	2021-07-26	CINECA, AVIO, ECMWF, SINTEF, IT4I, LINKS	Working version
1.2	2021-08-26	CINECA, AVIO, ECMWF, SINTEF, IT4I, LINKS	Main content completed
2.0	2021-08-27	CINECA, AVIO, MORFO, ECMWF, SINTEF, IT4I, LINKS	Pre-submission version, most comments resolved
2.1	2021-08-27	CINECA, AVIO, MORFO, ECMWF, SINTEF, IT4I, LINKS	Pre-submission version, all comments resolved, cosmetic improvements
3.0	2021-08-31	CINECA, AVIO, MORFO, ECMWF, SINTEF, IT4I, LINKS	Final version

Table of Contents

Glossary

Acronym	Explanation
AI	Artificial Intelligence
ANN	Artificial Neural Network
API	Application Programming Interface
BLAS	Basic Linear Algebra Subprograms
BPMN	Business Process Model and Notation
CFD	Computational Fluid Dynamics

CHT	Conjugate Heat Transfer
CPU	Central Processing Unit
DAG	Directed Acyclic Graph
DOF	Degree of Freedom
DS	Design System
FDB	Field DataBase - domain-specific object-store for Meteorological and climatological data
GB / GiB	Gigabyte = 1000MB / Gibibyte = 1024 MiB
GPU	Graphics Processing Unit
GRIB	GRIdded Binary or General Regularly-distributed Information in Binary form
HDF	Hierarchical Data Format
HPC	High-Performance Computing
HPDA	High Performance Data Analytics
IB	Infiniband
ICON	Icosahedral Nonhydrostatic Weather and Climate Model
IFS	Integrated Forecasting System
IO	Input/Output
KB / KiB	kilobyte = 1000 Byte / Kibibyte = 1024 Byte
LES	Large Eddy Simulation
LPT	Low Pressure Turbine
MB / MiB	Megabyte = 1000KB / Mebibyte = 1024 KiB
MPI	Message Passing Interface
NetCDF	Network Common Data Form
NWP	Numerical Weather Prediction
PGEN	Product Generation (ECMWF post-processing)
RANS	Reynolds Averaged Navier Stokes
REST	Representational State Transfer
TB / TiB	Terabyte = 1000GB / Tebibyte = 1024 GiB
UDF	User Defined Function
URANS	Unsteady Reynolds Averaged Navier Stokes
WRF	Weather Research and Forecasting Model

WRF-DA	Data Assimilation module for WRF

List of figures

List of tables

Executive Summary

The **ACROSS** (HPC BIG DATA ARTIFICIAL INTELLIGENCE CROSS STACK PLATFORM TOWARDS EXASCALE) project will build **an exascale-ready, HPC and data-driven execution platform**, supporting modern **complex workflows mixing HPC, BD and AI high-level tasks**, by leveraging on an **innovative software environment** running upon advanced **heterogeneous infrastructural components** (including GPUs, FPGAs and neuromorphic processors), as well as innovative smart resource allocation policies and job scheduling algorithms, up to the management of tasks inside jobs (pipelines, DAGs).

As such, envisioned ACROSS platform will provide high performance, and maximize resource utilization and **energy efficiency**. Also, through this execution platform, ACROSS project will demonstrate **value creation and innovation generation** for aeronautics, weather and climate forecasting (supporting for smart farming and innovative water management), technical and energy and carbon sequestration industrial sectors, as well as it will show the potential of the new EuroHPC supercomputers' ecosystem.

Objectives of the deliverables:

Deliverable 2.1 is related to Milestone 1 (Awareness of project objectives and requirements, M6) with the objective to give technical requirements and constraint as input for the Deliverable D2.2 (Description of Key technologies and Platform Design) the information contained in this deliverable will be exploited during the co-design phase towards Milestone 2 (ACROSS Key technological and Platform specification, M9) where we will put in place the overall ACROSS architectures co-design with technologies, platform and software choices.

Main Objectives of D2.1 are the following:

- Pilots identification and description (Pilot questionnaire)
  - Pilot characteristics
  - Pilots/Use case scenarios and applications functional description
  - Workflow detailed description, representation and complexity identification
  - Hardware / software baseline identification
- Pilots requirements for improvement (Improvement plan and requirements documents)
  - use cases improvement plan identification
  - Expected KPI
  - Current and expected performance
  - Improved workflow description
  - Co-design requirement

Position of the deliverable in the whole project context

Deliverable D2.1 is linked to WP2 "Cross stack convergence & co-design for HPC and Data driven HPDA software environment" dedicated to the co-design, key technology and platform identification, infrastructure set-up and integration and lesson learned. The D2.1 is in relation to the Milestone 1 "Awareness of project objectives and requirements" with due date at M6.
The D2.1 cover the M1-M6 co-design activities related to Pilots requirements.

WP2 is focused on co-design and pilots requirements and technical specification (see Figure 1), it will receive the baseline and improved requirements from Pilot WPs (WP5, WP6, WP7), support related to multilevel orchestration from WP4 and support for hardware and acceleration from WP3.

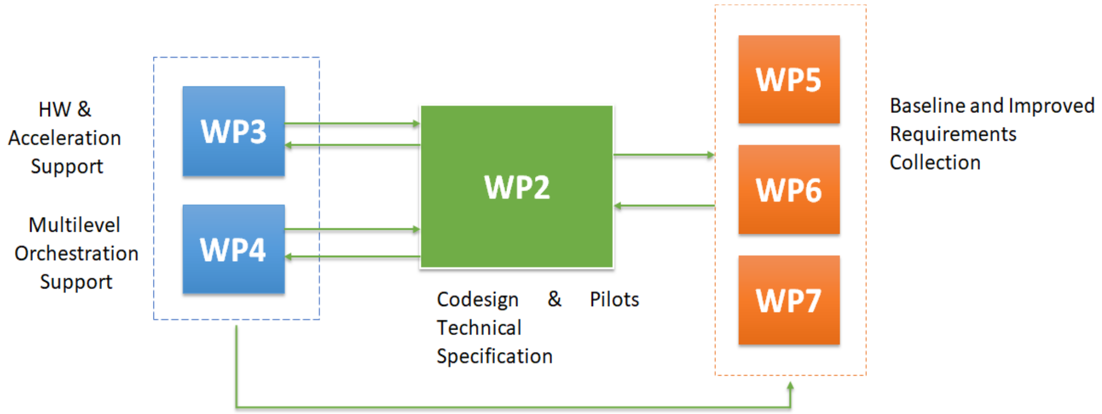

Figure 1 WP2 position in ACROSS project

Description of the deliverable

This deliverable reports the summary of the Pilots co-design requirements as they emerged from a collection of WP5, WP6 and WP7 "questionnaires", filled by all interested partners.

The document also describes the actual status of each Pilot (and of each relevant use-cases) with respect to a series of useful quantitative and qualitative indicators and the expected outcomes and improvements for that Pilot. Below are the most relevant co-design requirements for each Pilot based on the WP5, WP6 and WP7 integral "questionnaires":

- WP5 - Greener aero-engine modules optimization Pilot Requirements: The main objectives set by GE Avio Aero are to reduce the productivity target (reduction of time-to-design with respect to the current situation) by at least 50%, acting on both the modelling aspects of the physical problem and the optimization of the hardware. The other aspect is to improve the quality of the numerical results, getting as close as possible to the experimental reference results. Improvements of several aspects of the overall procedure (elimination of I/O issues and waiting times), new workflows. Moreover, the target of this activity is to develop an innovative, data-driven, AI-powered DS for turbines, capable of switching from standard aero components to innovative, additive-enabled aero components, aimed at improving the efficiency of the low-pressure engines module.

- WP6 - Weather, Climate, Hydrological and Farming Pilot Requirements:
    - Demonstrating Numerical Weather Predictions workflow with IFS model resolution improved from the current 9.0km operational resolution to 5.0km.
    - Improving WFLOW (hydrological application) runtime performance when compared to today's capabilities at least by a factor 5, to enable full ensemble simulations.
    - Demonstrating hydrological simulations over Rhine and Meuse basins (220.000 km^2 area) adopting 1.0km model resolution.
    - Carefully tuning the configuration of multi-layer data stores available on ACROSS computing resources, to efficiently support data management requirements.

- WP7 - Energy and Carbon Sequestration Pilot Requirements:
    - Improving OPM Flow runtime performance scaling when compared to today's parallel capabilities, scaling to 1000 processes with reasonable efficiency.
    - Carrying out flow simulations on large grids for long-term migration scenarios (> 1000 years), on models with up to 100M cells.

- Running direct flow simulation on models consisting solely of processed seismic data, at high resolution, with automatic and dynamic coarsening/refinement, on models with up to 100M cells.
- Demonstrating analysis of simulation results in-situ using methods from the AI spectrum in 3 new workflows.
- Increasing by 50% the overall data processing throughput (i.e., the number of scenarios evaluated per unit of time and the requests per second served in extreme cases).

It is important to remark that the pilot requirements, while challenging and comprehensive, do not necessarily fill the entire spectrum of HPC methodologies and technologies: in this sense the project scope will not be limited to satisfying the requirements presented in this document, but it will encompass a broader view of present and upcoming technologies for exascale HPC to collect best practices and to build a strong ecosystem.

1 Introduction

The Deliverable 2.1 summarizes the description of the pilots' co-design requirements for each Work Packages (WP5-7) in the framework of the ACROSS project, under the supervision of WP2. Specifically, the WP2 focus on the co-design effort in order to provide a simple access to ACROSS platform together with service provisioning for pilots (WP5, WP6 and WP7 respectively related to the Aeronautics, Weather and Climate, Energy and Carbon sequestration sectors) for the development of complex workflows on top of the HPC infrastructure.

The structure of Deliverable 2.1 is organized in following way:

- The Scope and Working Methodology related to the deliverable will be described, respectively, in Sect.1.1 and 1.2.
- The Overall Identification of the Pilots will be summarized in Sect. 2. The Sect. 3, 4 and 5 will be devoted, respectively, to the Pilot Requirements of WP5, WP6 and WP7. For each WP in these sections, the following will be reported:
  - the Expected KPIs;
  - the Current Workflow Performances;
  - the Proposed Improvements and Expected Performances will be detailed together with the explanation of the Improved Workflow;
  - hardware and Software Requirements;
  - the Data Management plan and the final Co-design Requirements.
- Sect. 6 will resume the findings for each pilot WP.

1.1 Scope

The scope of the Deliverable 2.1 aims to clarify the co-design pilot requirements inside the ACROSS project. This initial step is crucial to define innovative complex workflows based on advanced heterogeneous computation, innovative software and smart data management.
Specifically, well planned co-design requirements will drive new strategies devoted to improving the efficiency of the workflow execution, the Hardware/Software optimization, and the portability among different complex HPC architectures, in order to guarantee the compliance to the transition to pre-exascale and exascale modern HPC systems.
This step will ensure the ACROSS platform to handle optimally the orchestrated high-level tasks consisting on scheduling jobs and distributing the correct workload among the most appropriate hardware elements.

1.2 Related documents

ID	Title	Reference	Version	Date
D7.1	WP7 – Stage 1 requirements for HW/SW integration		v2.0	2021-08-31

1.3 Methodology

The devised methodology to ensure proper co-design approach for the creation of ACROSS platform consist of following activities:

- Joint WP2, WP3, WP4 audio calls
- Creation of the WP2 roadmap
- Pilot's use-case requirements, collected by means of the following documents
  - Pilot Questionnaires
  - Questionnaire Q&As
  - Improvement plan and requirements

Joint WP2, WP3, WP4 audio calls

The mutual decision of all partners was taken at the start of the project to host joint WP2 (Cross stack convergence & Co-Design for HPC and Data driven HPDA software environment), WP3 (Heterogeneous Hardware & Acceleration Support) and WP4 (Multi-level Orchestrator Towards Heterogeneous Exascale Computing) weekly audio calls for the duration of the first nine months of the project. This decision ensured that all important partners are present during the technical WP meetings.

For the first three months these joint audio calls have been led by WP2. Starting at M4 the decision was made to split the telco into three separate meetings – WP2, WP3, WP4 with specific agendas and lead by appropriate WP leader but to keep it still as a one joint meeting.

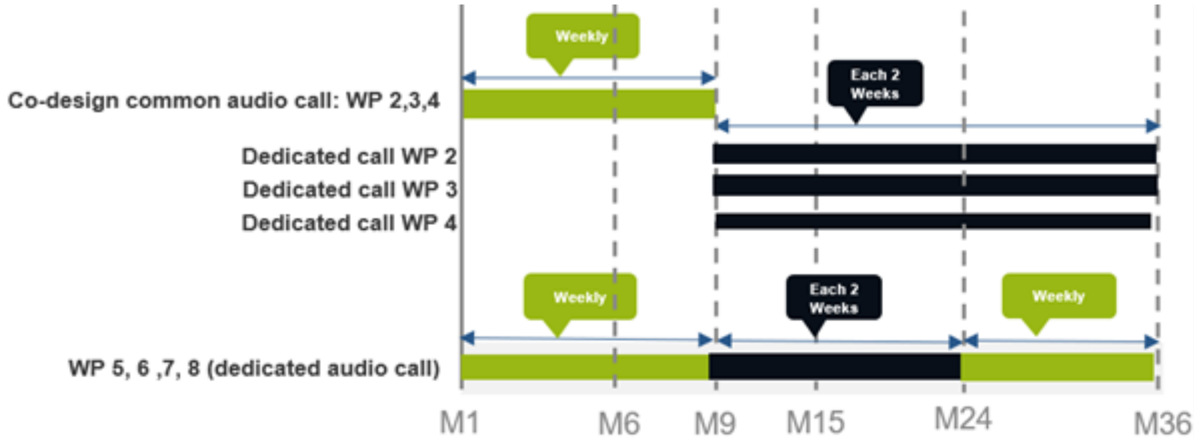

Figure 2 Joint WP2, WP3, WP4 audio call plan

Creation of the WP2 roadmap

To ensure an up-to-date overview of the planned activities related to the co-design tasks the WP2 roadmap was created. The roadmap contains all important activities such as planned pilot's use-case presentations, technical presentations, important deadlines, and milestones related to the co-design work package. The roadmap is being continuously updated based on the discussions and agreement from all partners.

Month	March		April					May				June			
Date	18.	25.	1.	8.	15.	22.	29.	6.	13.	20.	27.	3.	10.	17.	24.
Project's week	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Pilot's questionnaire	WP2+WP3+WP4 will create a questionnaire for high-level pilot requirements		Questionnaire presentation	Questionnaire open discussions				Questionnaire overall status presentation (only technical partners from WP2, WP3, WP4)		Pilot requirements recapitulation		2. Questionnaire overall status presentation (only technical partners from WP2, WP3, WP4)			
												Monday - Requirements status - WP5, WP6, WP7 (final)		TOC for the Improvement plan doc (to support D2.2)	First content for the Improvement plan
Technical presentations		CINECA and IT4Innovations infrastructure presentations	LINKS, ATOS - technology & infrastructure presentations	ATOS - FastML presentation			LINKS - Different computing technologies presentation								
		WP4 Orchestration approach presentation (LINKS, ATOS)			Low level orchestration presentations (Hyperloom, HEAppE)	Low level orchestration presentations (Damaris)									
Pilot's presentations						WP5 pilot presentation - current state				WP5 pilot discussion - update for ACROSS					
						WP6 pilot presentation - current state	WP7 pilot presentation - current state			WP6 pilot discussion - update for ACROSS	WP7 pilot discussion - update for ACROSS				
Deliverables															

Figure 3 WP2 Roadmap part 1

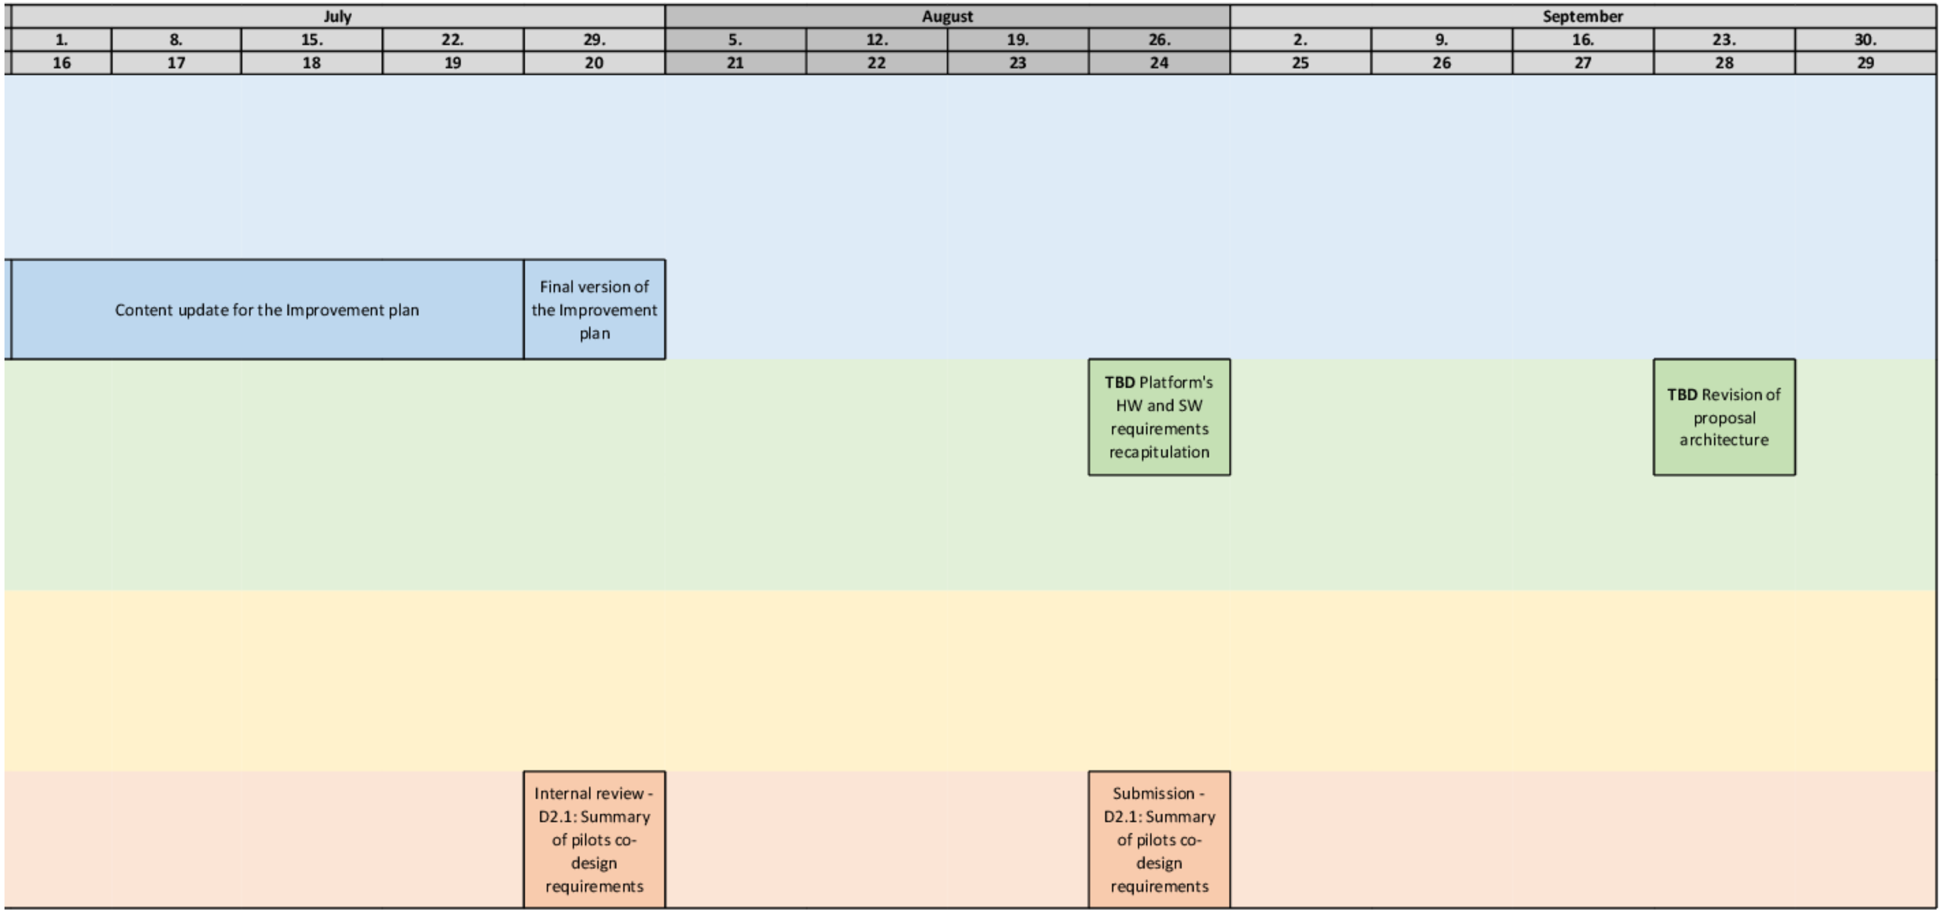

Figure 4 WP2 Roadmap part 2

Pilot's use-case requirements

Pilot questionnaires

To gather the pilot's requirements that will serve as a base for the co-design activities the 'ACROSS Pilot Questionnaire' document template was created. This questionnaire integrates questions from all technical WPs (WP2, WP3, WP4) and was provided to the pilot WPs (WP5, WP6, WP7) to fill in their answers.

Structure of the Pilot questionnaire:

- Identification of the pilot
    - *Owner of the Pilot*
    - *Title*
      Generic Name of the Pilot (or its Application Domain)
    - *Description*
      Functional description of the application in its operating context as well as its place in the lifecycle
    - *Usage scenarios*
      Examples to illustrate exploitation scenarios
    - *Actors/Systems*
      Roles of different actors/agents but also equipment (data centers, monitoring station, workstation, sensors, etc.) in the application
- Hardware and Software
    - *What software is used in the workflow?*
      Maturity: prototype, industrial exploitation, commercialized, etc.
      Sources: proprietary development, open source, etc.
      Parallelization and Programming Models: OpenMP, MPI, OpenCL, CUDA, etc.
      Open Source, Libraries, Independent Software Vendor products, etc.
      OS / programming language / Compilers
    - *What hardware is used by the workflow?*
      Computing technologies: CPU, GPU, FPGA, DSP, etc.
      Interconnection fabric: IB, Ethernet, etc.
      Storage devices (type + capacity)
      Data volume to be manipulated in memory and/or storage
- Workflow
    - *Provide the detailed description of the workflow.*
      Task graphs; Exploitation mode/constraints: application methodologies (e.g. real time, batch, periodicity, etc.); Information: (non-) structured, distributed data; Sources of data: dependence on other applications, etc.; Existing data base, data flow; Accessibility, confidentiality, etc.
    - *In terms of resource allocation planning, when do you expect your workflow should be ready for the deployment?*
    - *Is there a sequence to your workflow? Is there a way of knowing how you are progressing in this workflow?*
      Task dependency, decision-making structures (if then), specific constraints etc.
    - *Can you clearly split your pilot's workflow into individual jobs? If so, how many jobs are there?*
    - *Can the application workflow be split (decomposed) in sub-workflows that are executed on different computing systems (HPC, Cloud or AI-dedicated facility)?*
      For instance, an application workflow can be decomposed in one sub-workflow that requires to run on HPC cluster and another one that need to run on Cloud resources (e.g., to get access to dedicated accelerators –FPGA, etc.). Do you see a case such that in your pilot application?
    - *Scalability (strong/weak) of the pilot application: is there any scalability graph related to the whole pilot application or to any part of the application?*
      Is the pilot application already scalable on HPC supercomputers? If yes, how many nodes does the current implementation of the application use?
    - *Considering the case (if any) where your application workflow (or a sub-workflow –see question #11) runs on Cloud resources. In this case, which is the environment setup needed?*

Does the application require VMs and/or containers? If yes, which is the VM type needed (OS type, number of resources required, etc.)? And which is the containerization technology used (which is the technology used e.g., Docker, Singularity, etc.)?

- *What AI techniques are used in your workflow?*
  Nature (Symbolic vs Connectionism), ML, data sets, (federated) supervised/unsupervised training, inference constraints, etc. What frameworks do you use (TF2, Pytorch, etc.)?
- *Do you have any specific monitoring requirements?*
  Infrastructure level monitoring - utilization - Used CUPs, GPUs, data movement
  Application level monitoring - current status of WF monitoring - stdout like...
- *Does your workflow require user interaction at any stage?*
- *What visualization requirements do you have. Do you require access to a graphical interface (e.g. VNC/X session)?*
- *Have you used workflow languages to describe your problem (e.g. CWL - https://www.commonwl.org/)? If not, would you consider using a workflow language for your job?*

- Data management
  - *How do you define the size of your job? Where is this determined, and can it/does it change dynamically depending on the processing required?*
    Size as in the scope of the volume of processed data, requested/used resources - cpu time, gpu time, memory utilization, etc.
  - *Will you need to import data into the platform? If so, how will you import this data? How frequently will you need to import new data sets? What size are these data sets?*
  - *Do the workflows use different type of storage systems? Do datasets need to be encrypted or compressed? Are workflows sensitive to data movement latencies?*
- State-of-the-Art questions and how ACROSS will help
  - *What is the reasoning behind your current solution?*
  - *Are there any constraints regarding your current solution?*
  - *Are there any alternatives to your solution?*
  - *What is your vision beyond State-of-the-Art for your solution improved by ACROSS technologies?*
  - *References to existing tools/systems as well as to work in progress in the same domain.*

<u>Questionnaire Q&As</u>

The fulfilment of the Pilot's questionnaires was an iterative process. To speed up the interaction between the technical partners asking questions and pilot partners answering them the 'Questionnaire Q&As' document was created. Technical partners asked specific questions or requested more concrete answers regarding the Pilot's answers in the questionnaire. Pilot partners provided answers directly to the Q&As document and based on the mutual discussion also updated the answers within the Questionnaire itself.

<u>Improvement plan and requirements</u>

As the Pilot questionnaire document is focused mainly on the current state of the pilot's workflows, the 'Improvement plan and requirements' document was created to describe the planned improvements in more detail and to specify the requirements needed to achieve them. This improvement plan is filled based on the mutual collaboration and concrete discussions of WP3, WP4 and pilot's WPs. The resulting improvement plans will serve as a baseline for the co-design activities to support the pilots in achieving these improvements.

Structure of the Improvement plan and requirements document:

- Pilot improvement #No.
  - Baseline
    Reference to the baseline workflow item to be improved (e.g., LES CFD, workflow execution, accelerating a ML/DL model on NNPs, etc.)
  - Technological domain
    Improvement field of competence (e.g. data movement, code acceleration, workflow orchestration, etc.) and related tech WPs
  - Improvement Rationale
    Why the improvement is needed with respect to the baseline workflow?
  - Description
    List of functional description of the proposed improvements

- Expected results & KPIs
  Expected (quantitative where possible) improvement over baseline and definition of relative KPI (e.g. Time to Solution speedup in percentage, etc.)
- Improved workflow BPMN
  Description of the new workflow's scheme, if the improvement will determine a modification of the baseline scheme. In this case, please highlight which parts are changed or which parts imply a change in the non-functional HW/SW requirements.
- Technical requirements: Hardware
  List of improvement's hardware requirements (e.g. FPGA, NNP, etc.) and description (if needed)
- Technical requirements: Software & Orchestration
  List of improvement's software requirements (e.g. HyperTools, TensorFlow, etc.) and description (if needed)
- Feasibility degree
  How much effort and time will the improvement's implementation require by the partners? Please provides a feasibility category for the improvement (committed, hard, easy) followed by some additional observations/notes.
- Quality of end product
  - Standardization/Portability
  - Competitivity (vs SoA)
  - Evolution towards new computing technologies

1.4 Mapping with ACROSS Overall Objectives

In terms of requirements, the activities and co-design methodology set up for Pilot's linked to the Milestone one covering M1-M6 was focused to cover in part O1.1 ACROSS Overall objectives for co-design and O2.1, O2.2, O2.3 ACROSS Pilot specific objectives.

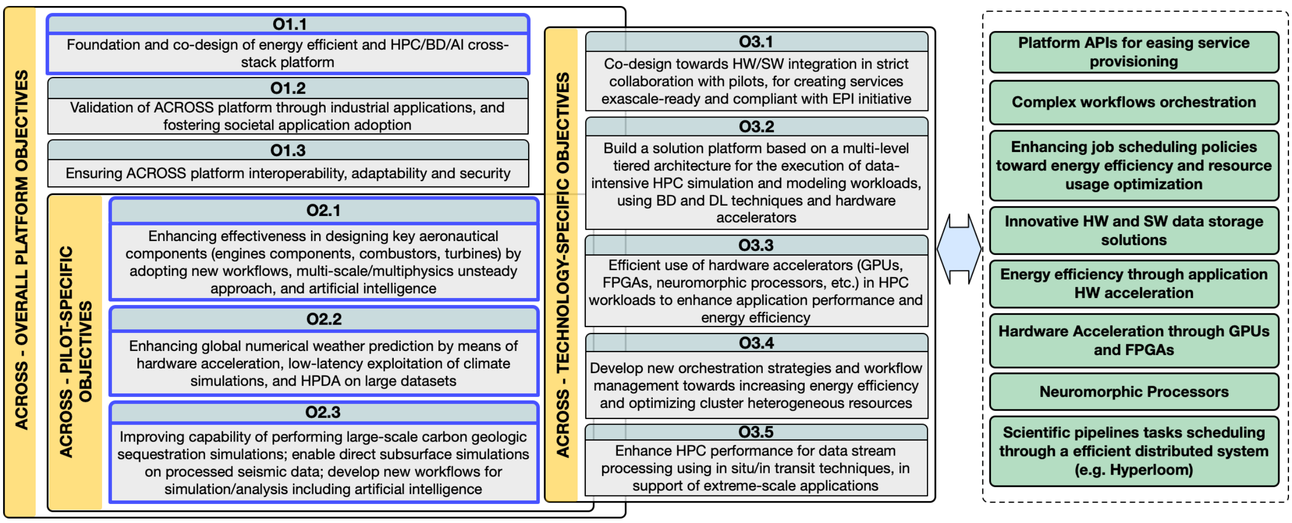

Figure 5 ACROSS Objectives

O1.1 Foundation and co-design of energy efficient and HPC/BD/AI cross-stack platform.

- Co-design activities for pilots' requirements

AERONAUTIC PILOT:

O2.1: enhancing effectiveness in designing key aeronautical components (engine components, combustors, turbine) for adopting new workflows, Multi-scale/Multi-physics unsteady approach, and artificial intelligence.

- Introducing new, complex design workflows for the design of key aeronautical engine components
- Innovating the existing combustor design process by introducing a Multi-scale/Multi-physics

- Developing an innovative, data-driven, AI-powered DS for turbines

WEATHER AND CLIMATE:

O2.2: enhancing global numerical weather prediction by means of hardware acceleration, low-latency exploitation of climate simulations, and enabling HPDA on large datasets.

- Improving the existing operational system for global numerical weather prediction, post-processing and data delivery, by exploiting hardware-acceleration and data streaming/object store techniques
- Enabling low-latency exploitation of climate simulations by integrating data delivery through a domain-specific object store;
- Developing and demonstrating an environment for user-defined in-situ data processing.

ENERGY AND CARBON SEQUESTRATION:

O2.3: Improving capability of performing large-scale carbon geologic sequestration simulations; enable direct subsurface flow simulations on processed seismic data; develop cross-stack workflows for subsurface simulations / analysis

- Enabling simulation of large-scale geologic carbon sequestration over very long timespans
- Improving existing and creating new subsurface applications by being able to perform direct flow simulation
- Developing new and improved workflows exploiting large-scale simulation capabilities and advances in AI techniques

2 Identification of the pilots

HPC systems and Data centric environments are use to solve complex workflows, intensive computing applications for large simulations and large amounts of data that require to be analyzed with the adoption of artificial intelligence techniques, ranging from scientific applications to engineering and industrial key sector applications. ACROSS project intend to demonstrate value creation and innovation in an Exascale ready, HPC and data driven execution platform for:

- Aeronautics (industrial key sector)
- Weather and climate forecast supporting smart farming and water management and Energy (scientific)
- Energy and Carbon sequestration industrial sectors (scientific)

2.1 Greener aero-engine modules optimization Pilot

Avio Aero has launched a challenging research activity aimed at significantly improving the feasibility and exploitation of advanced numerical modelling capabilities for critical engine components. The synergy between new generation HPC platforms, novel software(s), Innovative AI techniques and HPDA methods will open a new scenario for the design and optimization of aero-engines, providing unprecedented levels of accuracy and detail.
Avio Aero will leverage the state-of-the-art HPC resources available in ACROSS to verify the feasibility of these ambitious objectives. Two aeronautical engineering case studies will be rolled out: one regarding the combustor and the one referring to aeronautical Low-pressure turbines design. Both numerical investigations are based on complex CFD analyses and rely on CPU-intensive and time consuming routines. The objective is to demonstrate the speed-up opportunities given by state-of-the-art computing systems and to develop and deploy an efficient management of numerical results leading to extremely accurate performance prediction, never achieved before.

Turbine Workflow (Summary)

To generate the large database needed by the new DS developed in the ACROSS project, two kinds of CFD calculations will be carried out: URANS and LES. More in detail, once the design space will be defined in terms of the most relevant aerodynamic design parameters, it will be populated with a large number of optimal solutions. Each solution will result from a topology optimization carried out adopting URANS calculations. Once the design space will be populated with URANS solutions, a minimum number of points will be identified where high-fidelity LES calculations will be performed. The computational domain is defined on the basis of previous experience, and it consists of a multi-row, repeating stage environment.
The optimal URANS solution is analysed and obtained by means of “classical” ANNs that explore a limited number of degree of freedom (DOF) [1]. The LES, high-fidelity and high accuracy data are instead analysed by HPDA routine to understand the physics at hand. HPDA will be used to rapidly identify the mechanisms that generate the aerodynamic losses and hence a possible route to performance optimization. Finally, all the data collected by these classical procedures will be analyzed by different AI techniques, focusing on Bayesian regressions, ANN and advanced deep neural networks, in order to overcome the main limitations of the current design practice.

The Turbine-related workflow is shown in the Figure 6 below.

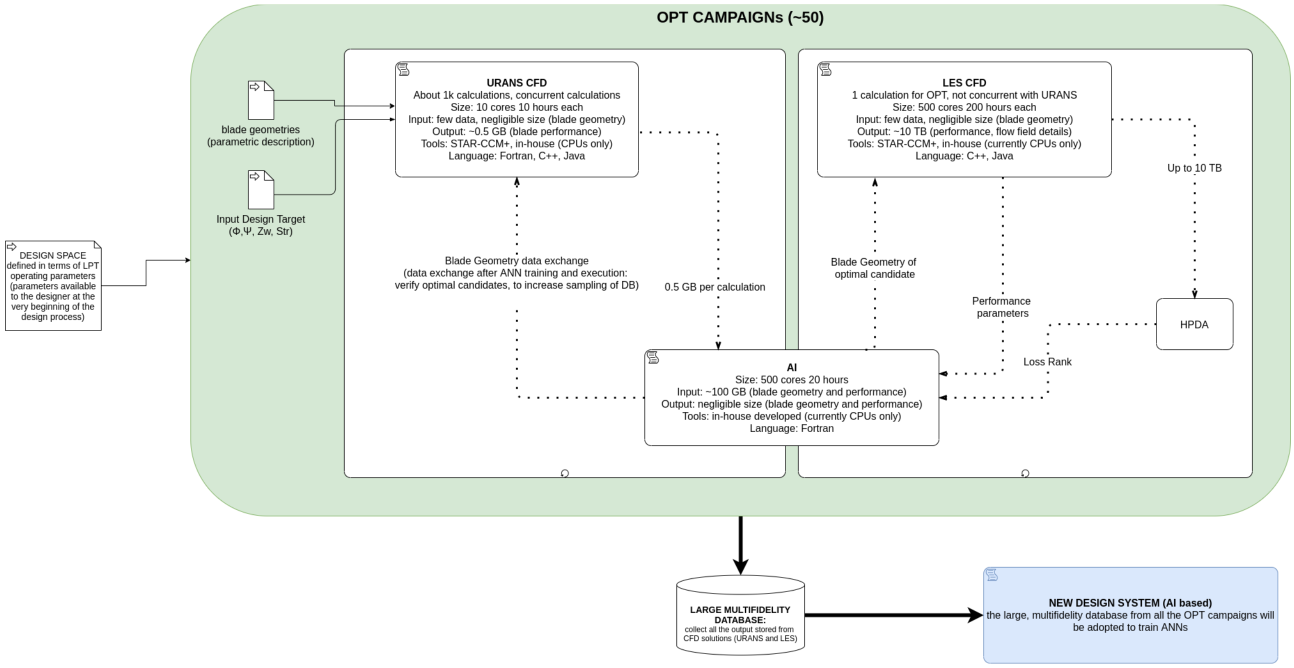

Figure 6 Turbine pilot Workflow

Finally, a scheme of the new DS is reported in Figure 6.

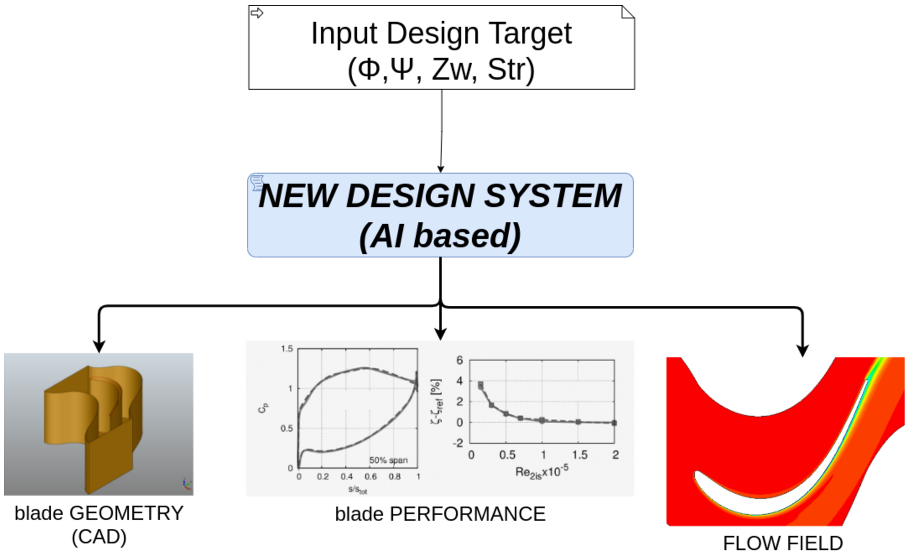

Figure 7 New turbine DS

Combustor Workflow (Summary)

The prediction of combustor wall temperature requires solving all the involved physics with a so-called multiphysics Conjugate Heat Transfer (CHT) approach, where different numerical domains are used to solve convective, conductive and radiative heat transfer modes. Moving to high-fidelity CFD methods to simulate the reactive fluid flow requires facing the problem in an unsteady fashion and thus having to deal with different characteristic time-scales of the involved heat transfer mechanisms. U-THERM3D is a CFD based multiphysics/multi-scale method where different domains are solved in parallel with a specific optimal time-step in order to limit the computational cost. U-THERM3D has been developed within the ANSYS Fluent code platform and consists of a set of User-Defined Functions (UDFs) and Scheme scripts which manage the unsteady domain integration, the parallel execution of runs on HPC platforms and the data exchange among solver processes. The basic idea behind U-THERM3D procedure is a de-synchronization of time-steps in the solution of the involved phenomena that can be summarized in convection (including several sub-phenomena as combustion, spray evolution etc.), conduction in the solid domain and radiation. Each of them is solved in a dedicated simulation, each with a proper advancing time-step resolution related to the characteristic time-scale of the considered heat transfer mode, running with a parallel coupling strategy. The tool has been extensively validated against different combustor configurations including that developed in the framework of the EU-funded projects LEMCOTEC [2] and FIRST [3].

The U-THERM3D procedure requires 3 calculations to be solved in parallel. Numerically, the most expensive one is the convective simulation in which turbulence and combustion are taken into account, requiring a very small-time discretization ($\approx 10^{-6}$ s) to properly catch the characteristic unsteadiness of the reactive flow. In order to ensure good solution quality, four flow through times must be simulated, two for flushing the solution and two for averaging the instantaneous quantities. An unsteady approach is also employed for the solid domain where conduction is solved with a longer time-step ($\approx 10^{-3}$ s) due to the longer characteristic time-scale of the conductive heat transfer mechanism. Since radiative heat transfer has a much smaller characteristic time-scale than the other two involved phenomena, it is solved with a steady-state approach. The numerical resources required for the latter two solvers are much lower than those allocated for the convective simulation due to the high number of solved governing equations and to the finer numerical grid generally employed for the gas phase calculation. However, these simulations have to run in parallel and, so, the time of the whole multiphysics CHT calculation will be fixed by the convective simulation which represents the bottleneck in terms of required computational resources. Once the set-up procedure of the three simulations is performed, the U-THERM3D procedure can start. Three independent Fluent jobs are launched at the same time and they will exchange data in fixed moments of the procedure, at the so-called "coupling time-steps" (usually after the solution of a fixed number of time-steps for each solver). If all the three simulations have not reached the same coupling time-step and exchanged data, the procedure does not continue and the calculations are frozen until the data exchange occurs. Usually, the convective domain (fluid/CFD solver) is the slowest as it must solve many equations (turbulence, combustion, spray evolution, etc.) with a very low time discretization within a fine numerical grid and so the radiative and conductive simulations are frozen until the data coming from the fluid simulation are received. After the data exchange and the re-synchronization occur, the simulations restart and continue until the next coupling time-step and data exchange event.

The complete workflow is shown in the Figure 8 below.

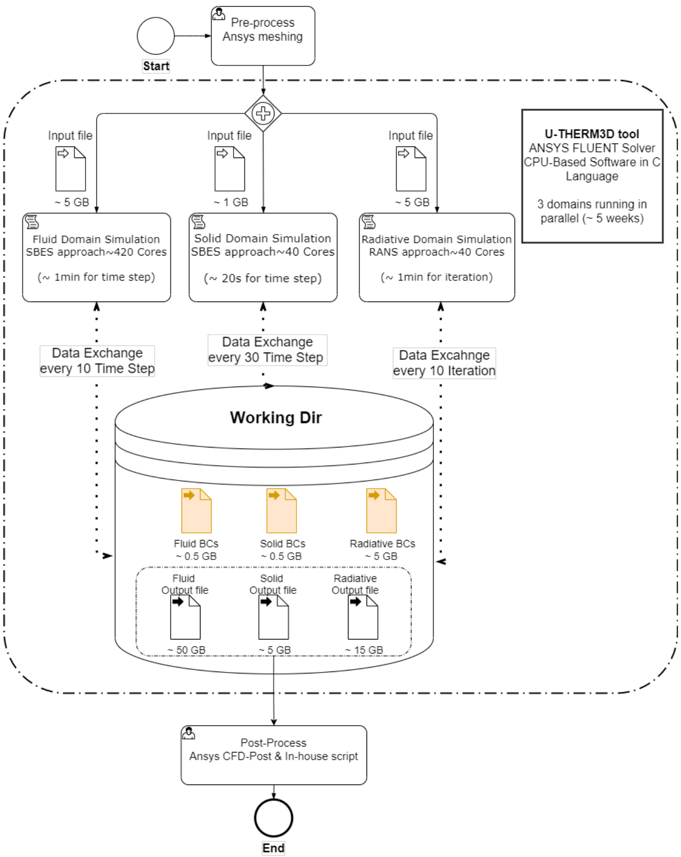

Figure 8 UTHERM3D tool workflow

2.2 Weather, Climate, Hydrological and Farming Pilot

The Weather, Climate, Hydrological and Farming Pilot aims to demonstrate the benefits of ACROSS infrastructure in the context of three deeply connected workflows. The main components of the workflows are:

- global scale numerical weather predictions based on IFS model and related processing jobs,
- climatological simulations performed with ICON global-scale model and ancillary applications,
- regional NWP downscaling performed with WRF and related data assimilation steps,
- hydrological simulations performed by WFLOW and
- farming services provided by NEUROPUBLIC.

The following diagrams summarize the interactions among such components, highlighting the central role of data management, based on the FDB domain-specific object store.

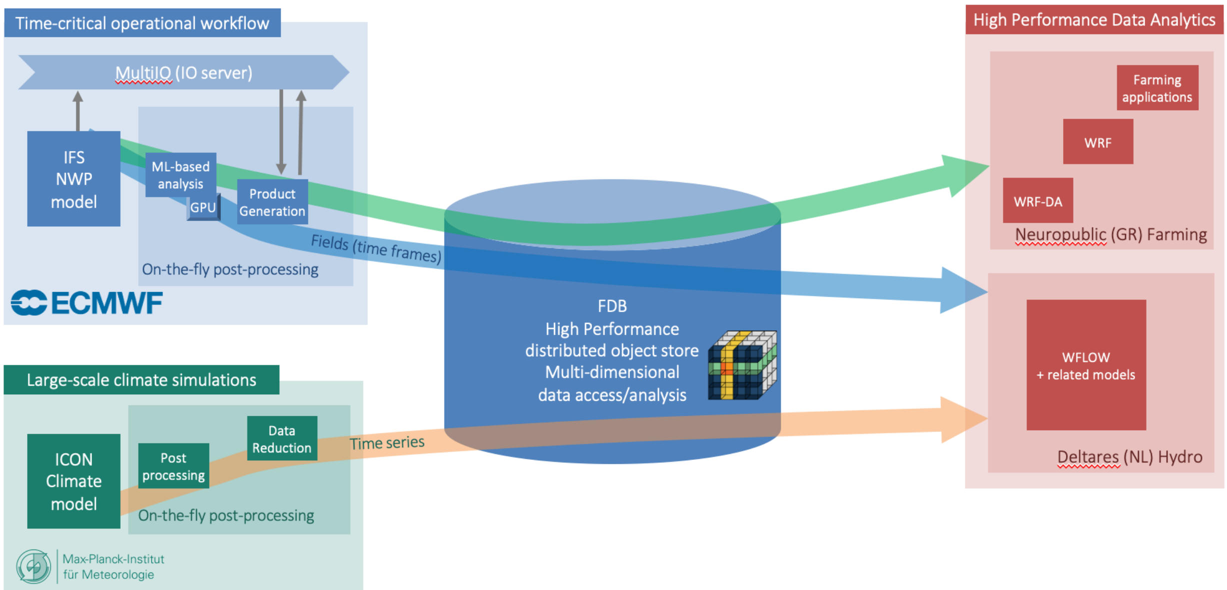

Figure 9 WP6 Workflow

More in details, we have the Hydro-meteorological workflow (blu arrow) composed by:

1. Global scale Numerical Weather Prediction workflow (ECMWF IFS model at 5km resolution + products generation)
2. Hydrological simulation over Rhine and Muse basins (Deltares WFLOW model at 1km resolution)

Hydro-climatological workflow (orange arrow) including:

1. Climatological simulation (ICON model – Climatological simulation at storm resolving resolution, 5 km)
2. Selection of a limited area domain
3. Hydrological simulation over Rhine and Meuse basins (Deltares WFLOW model at 1km resolution)

Farming advisory workflow. Global NWP, regional downscaling and farming specific post-processing applications (green arrow):

1. Global scale Numerical Weather Prediction workflow (ECMWF IFS model at 5km resolution + products generation)
2. ECMWF: on-the-fly post-processing (i.e. feature detection based on ML techniques)
3. Regional downscaling on Europe and Greece (WRF-DA + WRF-ARW up to 1km deterministic and 3km ensemble)
4. Farming-specific post-processing applications

Most of the numerical models (IFS, ICON, WRF) are already used in an operational context or are very mature (WFLOW), but their integration and the expected advancement are major results in the hydro-meteorological and climatological communities.

2.3 Energy and Carbon Sequestration Pilot

The reservoir simulator OPM Flow calculates future states for a porous reservoir connected to wells and possibly other facilities.

Since the known data are sparse, coming from seismic surveys and well data, one must typically use a large ensemble of realizations to account for uncertainty. One fundamentally important usage of a reservoir simulator is to provide the forward evaluation in a reverse-problem workflow where one seeks the true parameters (or at least sufficiently good values) of the underground reservoir. The ERT tool provides a tool for such reverse problems, by applying an ensemble smoother method, and can use OPM Flow for the forward evaluations.

Carbon sequestration use case
In the context of carbon storage, simulation is done to assess the safety and integrity of the storage site (avoiding leakage and fracturing of caprock) assessing the storage capacity and determining optimal well placements and schedules.

State-of-the art reservoir simulators have been development primarily for the oil and gas industry, where the focus has been on field-scale simulation models a few kilometres in extent and working on operational timescales up to 50 years. Such simulators have been used to study CO2 storage, but often research simulators with specialized models have been used, such as vertical equilibrium (VE) models that reduce the problem to a 2D problem. For simulating a regional-scale CO2-storage operation of the kind that could significantly contribute to climate change mitigation, the simulation model would have to be two to three orders of magnitudes larger than the typical field model. Moreover, simulations on sequestration timescales far into the future would be needed. If this is to be done in full 3D, computational requirements would far exceed what is currently used in the industry. The first use case aims to demonstrate such capability. This will require both excellent performance and parallel scaling, as well as specialized techniques that take advantage of the structure of the problem, for example exploiting VE in parts of the domain far from the injection wells, and reduced modeling (for example pressure only, or reduced resolution) in areas between injection fields.

Typical usage of the simulator software includes studying subsurface candidate sites for carbon storage, such as

- The Sleipner injection site.
- The storage site to be used by the Northern Lights project (Equinor, Shell, Total).
- The large-scale potential of the entire Utsira formation.
- Large-scale potential storage sites offshore Texas and Australia.

For a normal usage situation, an engineer or an engineering team sets up the workflow and interprets the results. However, much upstream work by geologists and others is required to produce the case.
Equipment used will typically be a workstation for small experiments, clusters or Cloud facilities for full workflows.

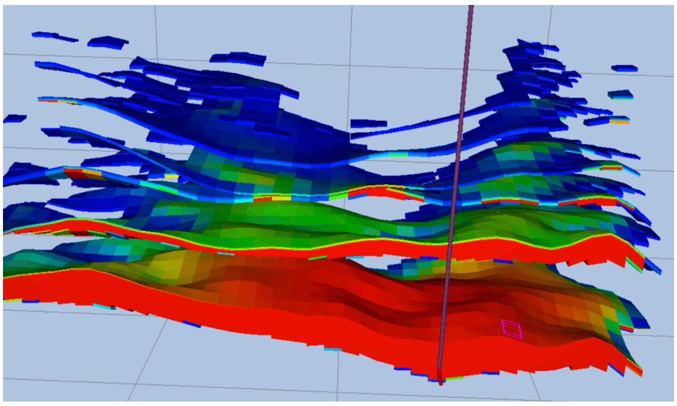

Figure 10 CO2 plume as simulated on the Sleipner benchmark case.

Seismic cube use case
New technology and workflows are increasingly making available large high-resolution seismic cube datasets with less human intervention than before. Reservoir simulation has typically been run on significantly upscaled, lower-resolution models that are created from such datasets through time consuming and work-intensive processes. To short-circuit those processes, ACROSS proposes to directly run simulations on the high-

resolution seismic cubes, thereby enabling a significant speedup of workflows. These cubes will have a much higher resolution than traditional simulation models, especially in the horizontal direction. The second use case aims to demonstrate that such high-resolution models can be simulated efficiently without sacrificing accuracy. This will require a highly efficient approach to dynamic coarsening and refinement.

3 WP5 - Greener aero-engine modules optimization Pilot Requirements

3.1 Combustor Design Use Case

To explore all possible ways of optimizing the tool, some possible academic test cases have been identified. In this first phase of the project, the TECFLAM combustor will be simulated. The TECFLAM swirl burner is a test rig developed from the cooperation between the universities of Darmstadt, Heidelberg, Karlsruhe and the DLR. The main goal of the experimental investigation is a detailed quantitative characterization of confined swirling-stabilized diffusion flames.
The presence of extremely detailed experimental measurements allows the complete validation of the numerical procedure. In addition, the combustor features an effusion-cooled plate, which represents the state-of-the-art in modern combustor cooling techniques, making this test case representative of the phenomena affecting industrial burners. Another important aspect is that the flame burns an air/methane mixture that is not bright and therefore it will not produce a relevant radiative heat load. This consideration allows to neglect radiation with a reasonable accuracy in terms of heat transfer characterization. Being able to work with only two domains (fluid and solid) will allow the U-THERM3D tool to be optimized with a low computational cost.

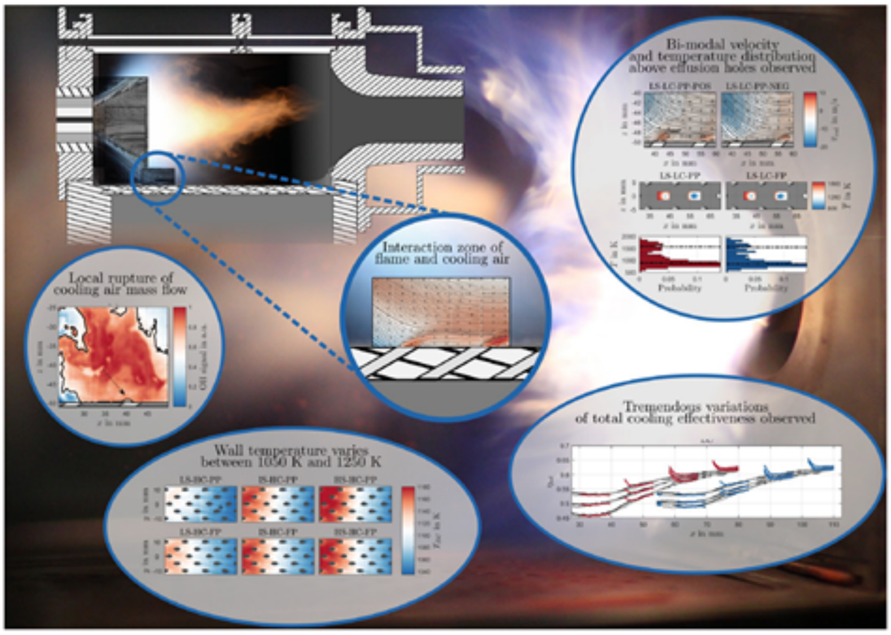

Figure 11 Graphical abstract with main measurements carried out on the TECFLAM combustor

The baseline workflow has been already reported at Sect.2.1, describing use cases related to Greener Aero-engine Modules Optimization pilot.

3.1.1 Expected KPIs

The following KPIs are defined for this pilot's use case:

KPI-1.1	Productivity target (time-to-design reduction with regards to current situation) for both aeronautical test cases. At least 50%, acting on both the modeling aspects and the optimization of the hardware
KPI-1.2	Combustor metal temperature prediction with regards to experiments. Target reduction of

	uncertainty margin to ±30K by acting on the modeling aspects used within the simulation

Table 1 WP5 Combustor KPIs

The main objectives set by the industrial team GE Avio are to improve the productivity target (reduction of U-THERM3D simulation time with respect to the current situation by at least 50%) acting on both modeling aspects of the physical problem and on the optimization of the hardware. The other aspect is to improve the quality of the numerical results, getting as close as possible to the experimental field results; in particular, the objective is to center the experimental temperature measurement with a margin of ±30K.This can be done by acting on the modeling aspects used within the simulation. As a final statement, expected results will be aimed at improving the quality of the numerical prediction with a reduction of the calculation time.
From this point of view, the objective is to achieve the same computing times as simulating only the fluid domain with adiabatic walls (theoretical limit that can be achieved by eliminating the waiting time for boundary conditions updating) but with the quality generated by a multi-scale/multiphysics approach as U-THERM3D.

3.1.2 Current Performance

Concerning the computational efficiency of the employed CFD solver (ANSYS Fluent), the following data relates to an aero-engine combustor that has 71 million of hex-core mesh elements and employs LES for turbulence modelling and a Probability Density Function (PDF) based approach as far as combustion modelling is concerned. Here, the liquid phase is also introduced by using a standard Lagrangian tracking. The performances are shown in terms of scalability, speed-up and efficiency up to more than 5000 CPUs, as reported below.

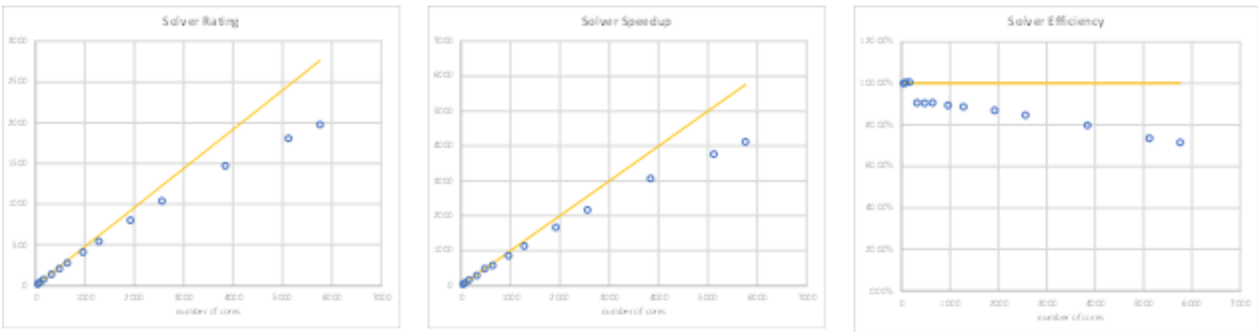

Figure 12 ANSYS Fluent rating, speedup and efficiency performance

A typical simulation based on the U-THERM3D tool of an industrial test case needs a computational effort of about 400 CPUh.

3.1.3 Expected Improvements

First of all, the fluid and solid domains (CFD and conductive simulations) will be solved in a single Fluent session using the “Solid time-step” feature that allows the two domains to be solved with a different advancing time-step. The tool will present a strong degree of customization since the principle of different boundary conditions for fluid-solid interfaces will continue to be used. In this new workflow, there will be no need to write/read profiles for the data exchange since it will be managed internally in the single Fluent session by proper UDFs (User Defined Functions). After the development of the new tool for solving fluid-solid domains, the solver for the radiative domain will be introduced. The idea is to still employ a different Fluent session for the radiative domain since, as already mentioned, the radiative heat load does not always impact significantly on the total heat flux. This makes it possible to keep the tool as flexible as possible to have the best performing version for each case study. In this sense, new radiative models will be analyzed to have a more accurate tool. Moreover, an optimization study will be carried out on the frequency with which the data exchange between the two remaining solvers should take place, based on both purely modeling aspects and hardware improvements (e.g., optimal balancing of computing resources, better workflow orchestration, etc.). All the changes described should lead to improved performance of the tool.

3.1.4 Expected Performance

The first improvement of the tool is related to the resolution of the fluid and solid domains in the same Fluent session. This new management will allow the elimination of I/O operations and waiting times, thus saving calculation time.

By introducing the radiative domain, an optimization of the frequency of data exchange with the fluid-solid solver must be carried out in order to minimise the waiting time.

During the ACROSS project, several aspects of numerical modelling will be considered in order to investigate which set-up guarantees the best prediction of metal wall temperatures of combustor. All these aspects will be combined with the improvement of workflow orchestration in order to guarantee an improvement in both software and hardware procedures.

3.1.5 Improved Workflow Description

Compared to the baseline, there are no longer three sessions resolved in parallel, but only two as shown in the Figure 13 below.

The new session consisting of the resolution of the CFD simulation and of the conduction within the solid has been represented with a single macro-block, but a data exchange between the two solvers has been maintained to highlight the fact that the idea behind the U-THERM3D tool has been kept unaltered. In fact, even if the improved procedure is based on the de-synchronization of the time-steps of the two domains, to use the most suitable set-up for both and therefore without increasing the computational effort; additionally, customized boundary conditions are used at fluid-solid interface to ensure a speed-up in achieving statistically steady-state conditions and convergence. The internalization of data exchange between the CFD and conductive simulations will be done via native functions of the ANSYS Fluent solver (UDF) and no longer via I/O operations that required simulations to be frozen. This will reduce the overall simulation time as there will be no more waiting time between simulations for updating boundary conditions and re-synchronization between different solvers. As can be seen from Figure 13 there are no major differences in the radiative domain compared to the baseline configuration. The reasons for this have already been partly indicated in the general section related to the description of the improvements. It is not always necessary to take the radiative heat transfer into account, so a configuration of this type allows for a more flexible tool. However, the possibility of internalizing this simulation as well will be evaluated during the ACROSS project. In any case, a study will be carried out to optimize the frequency of data exchange within the new procedure to harmonize calculation speed and the quality of the solution obtained.

Pre and post-processing operations are unchanged from the baseline configuration. Given the different geometries and test cases that can be simulated, a generalization would be very difficult to achieve. In addition, this type of operation is not carried out many times for a single simulation, so they are not considered to get a particular impact on the total time of the procedure.

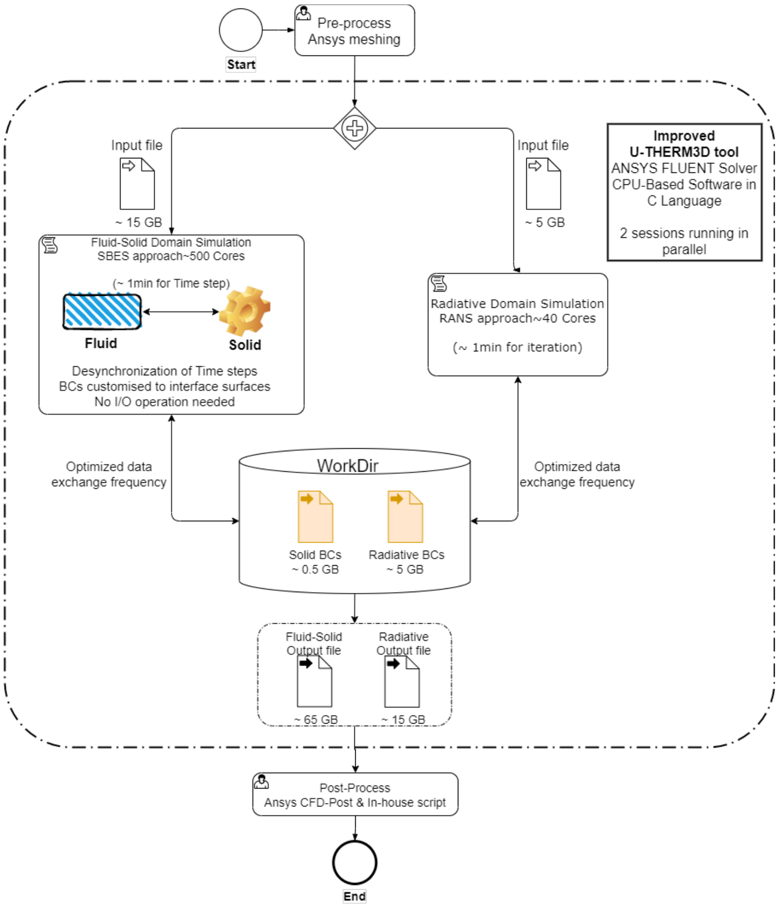

Figure 13 Improved workflow for U-THERM3D tool

3.1.6 Hardware Requirements

The adopted software is the commercial CFD solver ANSYS Fluent, is written in C programming language, run on X86-64 CPU architecture and requires MPI libraries to interact among the compute nodes via IB communication system

No additional hardware requirements, beyond those needed for the basic procedure, are expected for the optimized one. An analysis of the impact of the characteristics of the HPC architecture (compute node, memory size, interconnection topology) on the overall performance could however be carried out.

3.1.7 Software Requirements

ANSYS Fluent needs to use licenses that are allocated on the UNIFI servers, all the execution nodes of HPC cluster must be able to access software license (FlexLM) on an external server. Based on ANSYS Fluent, the Department of Industrial Engineering of University of Florence has developed the U-THERM3D tool for solving conjugate heat transfer calculations with a parallel loosely-coupled approach in the context of unsteady multi-physics problems.
To run the whole U-THERM3D procedure several simultaneous jobs must be launched. Hence, due to this need, a dedicated launch script is needed, capable of managing and splitting a single job into several ones.
The U-THERM3D tool has been preliminary validated against different combustor configurations, including the one developed for the LEMCOTEC European project and the EU-funded project FIRST; it has also been used for several industrial test cases owned by GE Avio.
ANSYS requires MPI libraries to interact between the HPC nodes. The U-THERM3D tool runs on Fluent version 2019R1 by means of User-Defined Functions (UDFs) and Scheme commands defined within the ANSYS Fluent framework. There are other software packages not directly used by the U-THERM3D tool but which are necessary for the generation of the calculation grid and for the post-processing of the results.

Specifically, ANSYS Meshing is employed to generate the mesh grid calculation whereas ANSYS CFD-Post is used for the post-processing phase.
In addition, in-house Python scripts are also adopted for more detailed post-processing analyses.
For the improved procedure, it may be necessary to install new versions of the ANSYS Fluent solver (2021R1 or later) able to use all the new potential features that will be implemented in the software and then proceed to their customization. For subsequent orchestration optimization, it may be necessary to employ specific software to monitor how the CPUs are working on the HPC cluster in order to understand their effectiveness and then proceed with an optimization.

3.1.8 Data management

The size of the data is strictly dependent on the type of test case studied and the type of simulation carried out. A typical simulation based on the U-THERM3D tool of an industrial test case needs a computational effort of about 400 CPUh. The selected test cases involve an output data quantity of about 120 GB. The results will be generated in the standard format of the ANSYS Fluent solver, i.e. a ".cas" file containing the calculation grid and the numerical setup of the simulation and a ".dat" file containing the complete simulation results.
After post-processing of the data, images, videos and data profiles can be obtained for dissemination of the results.

3.1.9 Co-design requirements

The new management of data exchange within the UTHERM3D tool will require a high effort from the WP5 team.
Improvement needs a completely new data exchange logic which in turn must be implemented within the ANSYS Fluent solver by in-house UDFs. After the implementation of the new workflow, a validation campaign must be carried out to ensure that newly developed software solutions respond well to customization and leads to a qualitative benefit in terms of results.
In terms of hardware improvements, the effort will be shared with the ACROSS partners.

Requirements summary:

- Baseline
  - One simulation
    - ~400 CPUh
    - 120 GB
- Improvements
  - New tool for fluid-solid domains - CFD and conductive simulations will be solved in a single Fluent session
  - New tool for radiative domain – different Fluent sessions for the radiative domain introducing flexibility for each use-case
  - Optimization study regarding the data exchange between the solvers – optimal balancing of computing resources, workflow orchestration, etc.
- HW requirements
  - CPU compute nodes
  - IB interconnection
- SW requirements
  - Commercial CDF solver ANSYS Fluent (C, MPI)
  - Access to UNIFI license servers
  - U-THERM3D tool
  - ANSYS Meshing
  - ANSYS CFD-Post
  - Python scripts
- Other requirements
  - Infrastructure monitoring tools (CPU utilization, etc.)
  - New management of data exchange within U-THERM3D
  - New data exchange logic implemented within ANSYS Fluent solver
  - Checkpointing

3.2 Turbine Use Case

The target of this activity is to develop an innovative, data-driven, AI-powered DS for turbines, capable of switching from standard aero components to innovative, additive-enabled aero components, aimed at improving the efficiency of the low-pressure engine module [4], [5]. The new DS will be based on a large numerical database developed within the project. Figure 14 below reports the main phases.

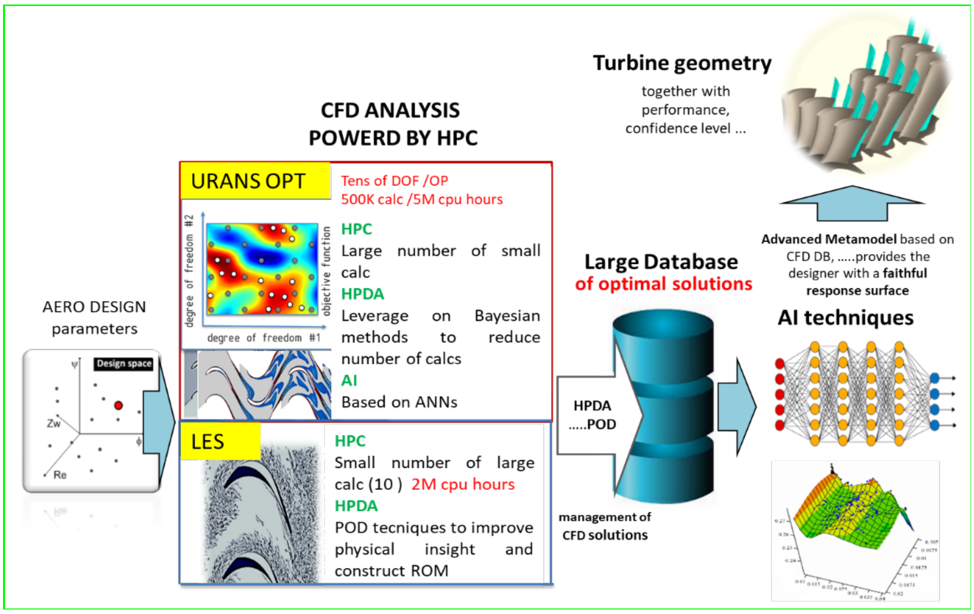

Figure 14: New DS for turbine blade profiles

Two kinds of CFD calculations will be carried out for this purpose: unsteady RANS (URANS) and Large Eddy Simulation (LES). More in detail, once the design space will be defined in terms of the most relevant aerodynamic design parameters, it will be populated with a large number of optimal solutions. Each solution will result from a topology optimization carried out adopting URANS calculations that will be further tested with LES. The results of LES will be further analyzed by means of HPDA and artificial intelligence (AI). By the conjunction of these procedures, we expect to build a fast DS for LPT that generates an optimal blade geometry for a wide range of design parameters and that is able to accurately predict the blade performance.

The baseline workflow has been already described in Sect. 2.1

3.2.1 Expected KPIs

The following KPIs are defined for this pilot's use case:

KPI-1.1	Productivity target (time-to-design reduction with regards to current situation) for both aeronautical test cases. At least 50%, acting on both the modeling aspects and the optimization of the hardware
KPI-1.2	Overall low-pressure turbine aero design efficiency gain with regards to current method, bringing reduction of specific fuel consumption (SFC). Target SFC reduction by 0.30% by using the innovative AI-driven DS

Table 2 WP5 Turbine KPIs

The expected KPI of this activity are related to the change of paradigm provided by the new DS. This new DS will be based on meta-models that will leverage on the generation of a large numerical database by means of HPC, and their analytic by means of HPDA and AI.
This will provide a new tool for designers that should be able to reduce by 50% the overall time to design an optimal 2D blade shape, with respect to current design approach based on long iterative procedures. The AI-

driven DS will also allow switching from standard to innovative aero components, thus improving efficiency and reducing engine weight. This is expected to bring a reduction of specific fuel consumption (SFC) of 0.3%.

3.2.2 Current Performance

To clarify the performance of the current approach, it is important to focus both on the design process as a whole and on the main time consuming tools adopted for the design.
As highlighted in Figure 15, the current workflow adopted for the design of a LPT blade, starts with a preliminary design. Simplified approaches (mainly correlation-based) are adopted in this phase to derive a first design of the blade, looking for a promising geometry for the specific design target. This preliminary design usually exploits very limited computational power, thus resulting in a not negligible wall time.
It is important to underline that preliminary design tools based on Company's expertise (coming from experimental and design data) are characterized by limited accuracy and typically do not allow obtaining the requested performance for the LPT blade. Furthermore, it is the industrial engineering experience that drives the final choice on the blade design, thus increasing the risk of errors.
After the preliminary phase, a successive optimization phase will be performed looking at an optimal solution for the specific target. Currently, the optimization campaign is based on CFD RANS calculations for performance prediction, and metamodel-based optimization techniques for the optimization process. This phase is usually very time consuming and requires relevant computational resources. Moreover, as a general approach is not defined, the process is usually an iterative one with a demanding contribution for the designer that has to drive the optimization. Once the optimum has been obtained, a final refinement is usually needed introducing some specific features or technologies that further improve the design and match all the constraints.
In the current approach, therefore, the whole process is repeated for each new design, as the experiences gained during the different campaigns have little or no impact at all on the successive activities.
The introduction of the new design system instead, will drastically change the overall approach to the design. Once that the new DS will be derived, it will be possible to obtain the optimal design directly inferring it from the large database adopted for training the ANN metamodel. As the new DS will leverage the data extracted from several OPT campaigns covering the whole design space for Aeronautical applications, it will be able to directly provide optimal solutions that yield the requested performance. Thanks to the systematic approach and the high accuracy of the calculations that stand behind the new DS, it is expected that it will be able not only to drastically reduce the time required for the design, but also to ensure a significant improvement in the LPT blade performance.

Focusing on the tools adopted during the optimization and within the ACROSS project, it is important to discuss the current performance of the numerical tools, as it will be improved during the project. These improvements indeed will play an important role in the feasibility and the cost/effectiveness of the overall process. No particular emphasis is placed on CFD calculations, as they are performed adopting a commercial code (STARCCM+) that currently runs on CPUs only, and no specific updates are expected during the project. However, it is important to underline that the code features very good scalability, thus resulting in an effective choice for exploiting modern HPC clusters.
URANS calculations are expected to run in around 10 hours per calculation using 10 cores, but as thousands of calculations are needed, several calculations must be performed simultaneously. Once the database of URANS solutions is produced, a metamodel approach (based on ANN) is adopted for each optimization.
The ANNs that analyze and find the optimal solution are very small since they manage a very low number of DOF (few tens). Each optimal solution will be finally evaluated by means of an accurate LES calculation that is run using around 500 cores with an expected wall clock time to about 10 days. The HPDA of LES results is currently performed after the end of the LES simulation, requiring up to 10TB of files saved in the work area. Typical running time of current HPDA routines on a dataset of 1.5TB is 10h on 500 cores [6].

3.2.3 Expected Improvements

It is important to underline that the introduction of the new DS into the Avio Aero industrial design procedure, will results in an important innovation and improvement of the traditional approach adopted for the turbine design. As highlighted in [6], the innovate DS, will be able to overcome the typical time consuming design approach (preliminary design + detailed optimization campaigns) with a single shot CFD driven, data centric one.

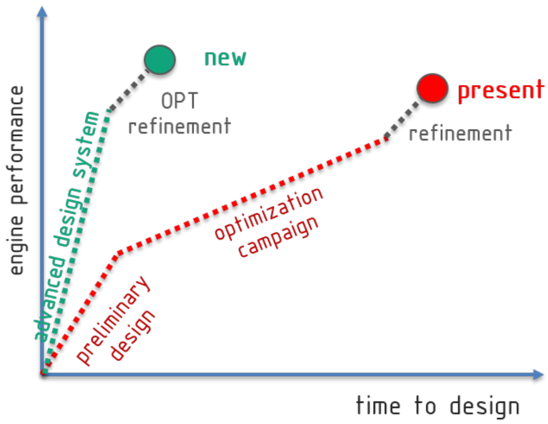

Figure 15: Improvements expected for the new DS

As already introduced in Sect. 2.1, to derive this new DS, the overall design space of interest for LPT in aeronautical applications will be defined. Each point of this design space mainly consists of different operating conditions, or different kinds of applications (i.e. engine size, vehicle range, etc.), that are typically encountered in the design activities. Within this design space, several points will be explored by means of specific optimization campaigns. These campaigns will be carried out adopting state-of-the-art CFD calculations (URANS and LES) and AI tools for optimizations. All the results obtained during these activities will lead to a very large database that will be adopted for deriving the new DS, leveraging on advanced AI tools. In the end, this will result in a system able to provide, in a very reduced time, the optimal geometry for every new design, thus replacing the very time consuming approach currently adopted and described in Sect. Current Performance.

In this context, however, the adoption and improvement of advanced tools in all the phases of the construction and exploitation of the new DS, will play an important role in making feasible, effective, and efficient the overall system. For these reasons, starting from the baseline workflow already presented in Sect. 2.1, two main fields for improvements have been identified, as highlighted in Figure 16:

1. The AI of the DS itself from hardware and software perspective
2. The integration of LES an HPDA from the orchestration perspective

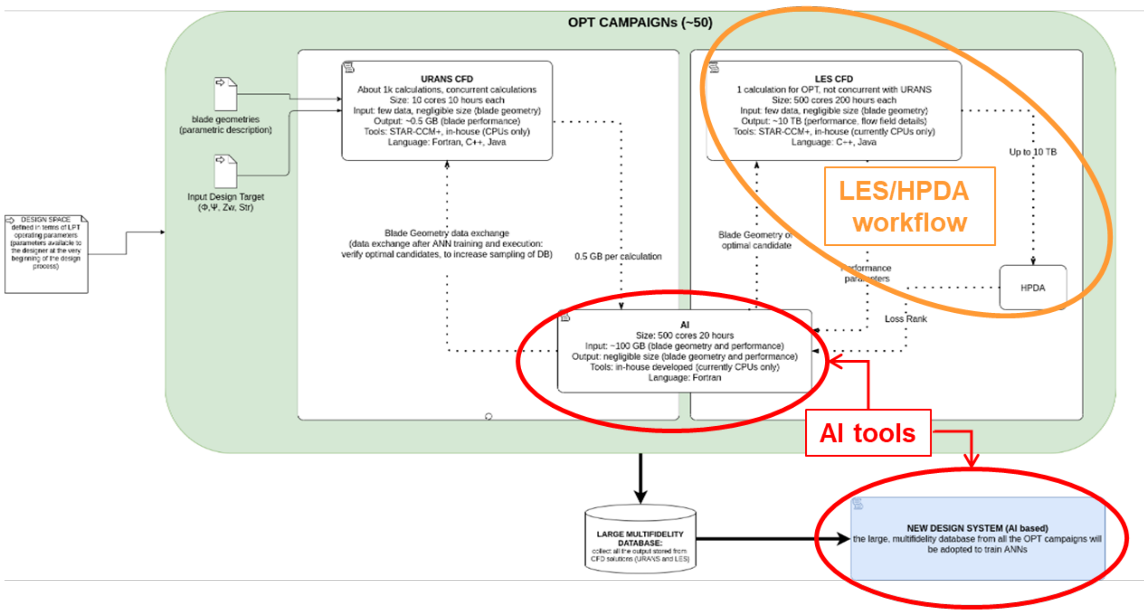

Figure 16: Baseline workflow and main fields for improvement

The new DS will need to manage the large database generated within the project with a large number of DOFs (up to thousands). The output will be multi-dimensional with the aim of providing the blade performances on the fly with a high accuracy by relying on the multi-fidelity simulations applied during the project. Therefore, this new tool will require large computational resources and effective languages and frameworks. State of the

art techniques and architectures will be exploited to speed up the whole system and innovative architectures will be also exploited to speed up the system. The improvement on the AI tools will be mainly focused on the adoption of accelerators to speed up the whole process (both for training and inference).
This part, therefore, will be mainly focused on tools porting to exploit accelerators and their tests. This will play a key role during the project, reducing the requested time in each optimization campaign, and, even more important, it will relevantly impact the effectiveness of the DS for daily use after the project, when it will be exploited for LPT blade design.

Concerning the second point, the introduction of HPDA while the LES calculations is still running may allow a fast evaluation of the hi-fidelity simulation convergence by saving time and stored data. In the current workflow, LES and HPDA are done separately and the HPDA is simply adopted as post-processing tool. The number of iterations of LES and the amount of saved data are determined a priori and chosen to support the statistical analysis, resulting in a large amount of data to be stored. The new perspective will adopt the HPDA to check the LES convergence on the fly on high statistical moments without impact on the LES computation speed. The HPDA (that is a fully data-driven analysis procedure) results will eventually be used to provide an efficient stopping strategy of the LES, that in turn will reduce the execution time by stopping the process earlier than what it is possible at the moment. From the orchestration viewpoint this will imply to devise an effective way to run LES and HPDA workflow tasks concurrently, by taking into account their specific requirements in terms of resources.

3.2.4 Expected Performance

It is important to underline that the KPI for WP5 (design time reduction of ~50%) is related to the introduction of the new DS as a whole, not to the performance improvement of the single components. As described in Sect. Current Performance, in the current approach, each LPT design consists of a lengthy and complex activity with relevant costs in terms of time and computational resources. By exploiting very large computational power, modern SW and HW frameworks, and leveraging on advanced AI techniques, during the ACROSS project a completely new approach will be derived, able to ensure an improvement in the performance of a new LPT blade and a drastic reduction in time required for its design.

However, improvements of the single components of the workflow adopted for deriving the new DS, as well as improvements in the workflow itself are still expected to contribute to the achievement of the KPI: for instance, the introduction of accelerators (GPUs, Neural Network Processors) in the AI phase is expected to provide an improvement of an order of magnitude on the AI section. In case of more exotic HW accelerators (FPGAs, neuromorphic processors), due to the effort needed to port code to those architectures, a more significant performance improvement should be demonstrated to justify their introduction in daily usage of the DS, hence their adoption is still under investigation.

For the improved workflow on LES/HPDA a time reduction at least of 30% is expected. Moreover, it is important to underline that it will also ensure a more homogeneous database in terms of convergence, as the same convergence will be checked and reached for all the calculations.

3.2.5 Improved Workflow Description

The improvement on the AI part will not alter the original workflow. A relevant change instead will be implemented in the HPDA/LES part.
The detailed workflow for LES and HPDA in the baseline procedure is reported in Figure 17. From this scheme it is clear that HPDA is performed downstream of the LES calculations running until reaching a fixed number of iterations.

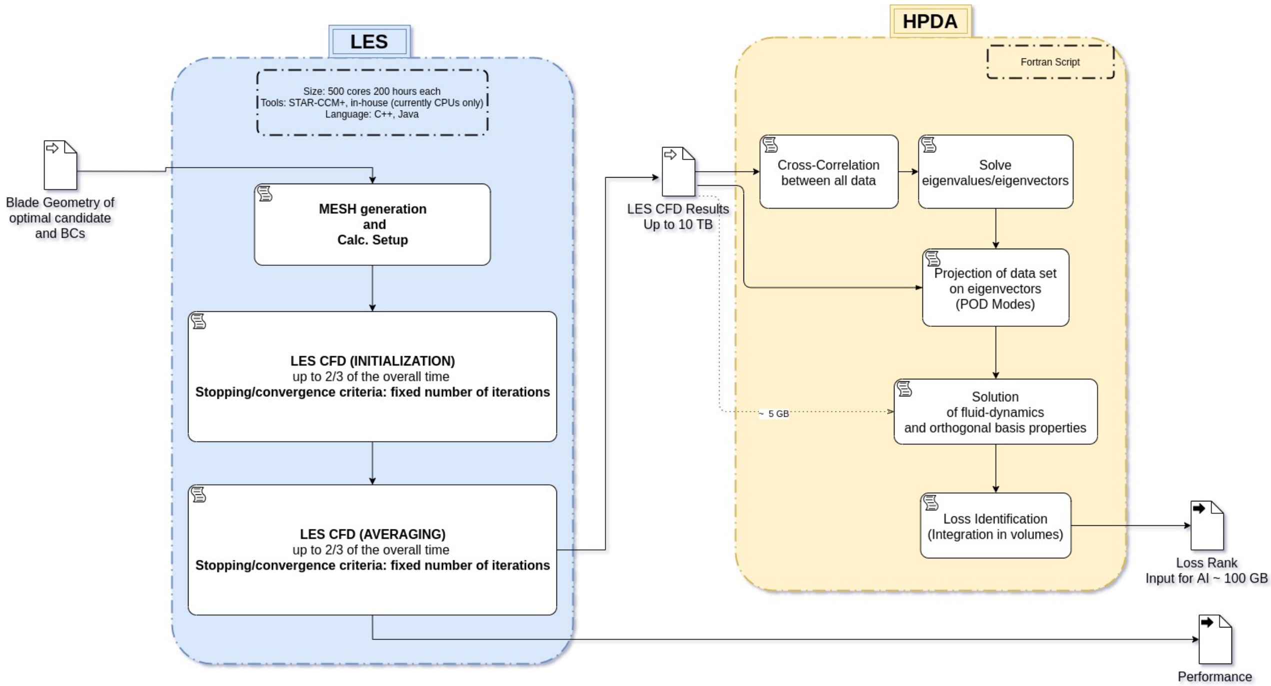

Figure 17: Baseline workflow for LES/HPDA

As discussed, in the improved workflow HPDA as well as some other performance evaluations will be carried out while LES is running, and the convergence of this co-processing will be adopted as a new, more effective, stopping criterion for the calculations. Moreover, all the results will be directly available at the end of the calculation, as post-processing is performed concurrently with the calculation itself.

The workflow of this new procedure is reported in Figure 18.

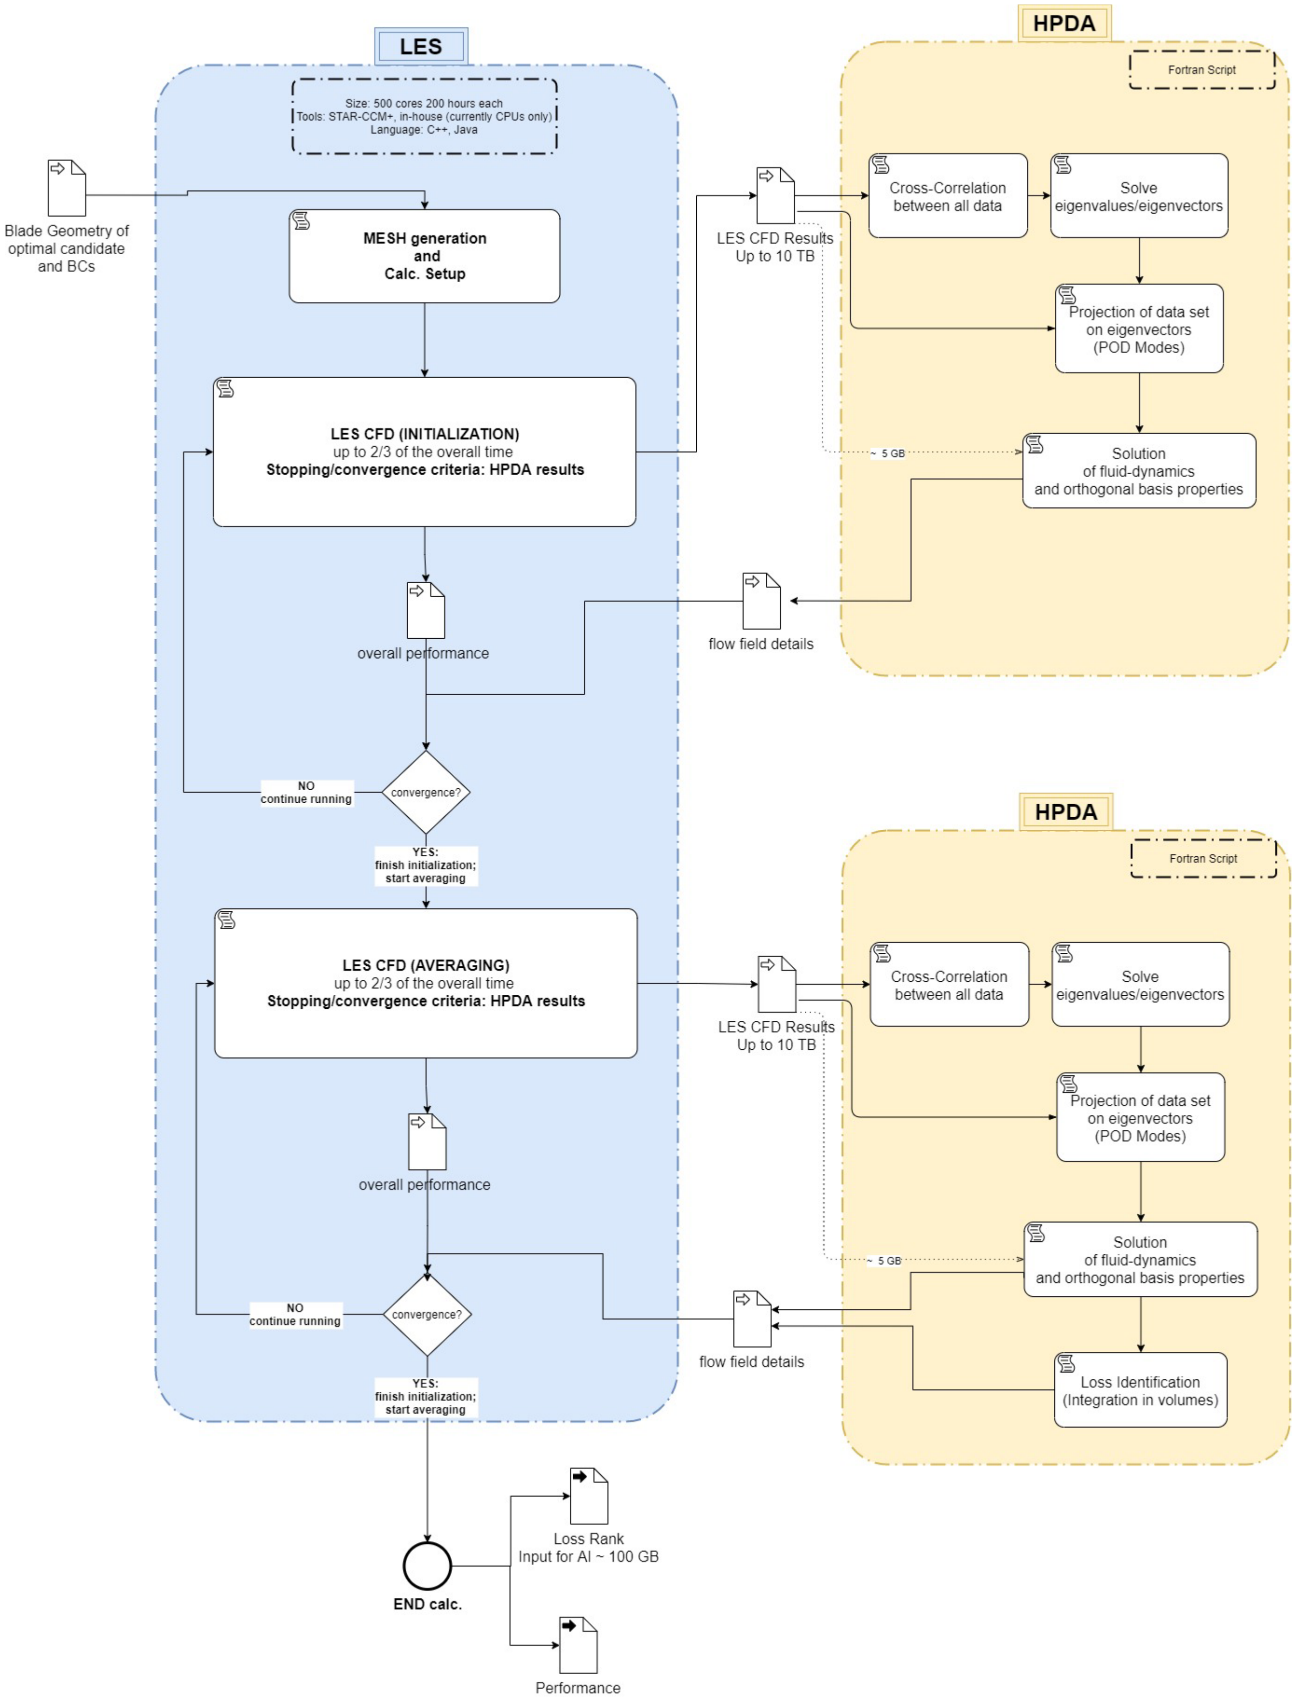

Figure 18: Improved workflow for LES/HPDA

3.2.6 Hardware Requirements

At the beginning of the project, all the tools adopted in the LPT pilot run on CPUs only. During the project, different architectures may be tested for further improvement. For example, the proposed improvements on the AI part, will mainly focused on the adoptions of accelerators. Starting from tools that run on CPUs only, the new advanced one should be able to exploit both CPUs and a range of AI-accelerators including general-purpose GPUs, FPGAs but also specific VPUs (Vision Processing Units) and NNPs (Neuron Network Processors) to test the potential benefit of innovative architectures.

No special requirements concern the improvements on the LES/HPDA part. Being the HPDA part of a set of open source routines in Matlab and Fortran90, it might be feasible to eventually test portability to GPU with help of WP3.

3.2.7 Software Requirements

CFD calculations (both URANS and LES) will be carried out using the commercial software STAR-CCM+. The very last release of STARCCM+ is needed: relevant updates have been introduced for turbomachinery applications. Morfo is currently using version 2021.1. STAR-CCM+ runs on CPUs only, therefore, for the moment we plan to use only CPUs for the project. However, a testing version running on GPUs (CUDA based) may be available during the period of the project. Finally, the only requirements for the AI improvements concern the availability of TensorFlow/Keras together with OpenVINO as inference interface to hardware, as the improved tools will adopt these frameworks, which, at the same time, converges towards the adoption of standard.

The software for the HPDA part is based on in-house developed routines with Matlab and Fortran. Namely, the Fortran90 routines are parallelized using MPI libraries and adopting LAPACK(MKL). In-house developed codes will be adopted for pre- and post-processing. These tools are mainly Fortran, Python, or Java applications. The availability of Paraview will be also necessary for post processing.

3.2.8 Data Management

Data will have a central importance in turbine-pilot as it is mainly focused on the development of a new DS based on a huge database built during the project and leveraging advanced CFD modelling, HPC resources and HPDA techniques. For these reasons, the data generated during the whole project will play a key role for reaching the objective of the pilot.

The articulated workflow of the current pilot suggests an equally complex data management during the whole procedure, with several points in which data are generated and exchanged among the different tools adopted. The whole process starts with few scalars as input. From these values, a blade geometry is defined and a first set of data therefore is generated in this preliminary phase to provide CFD calculations with the geometry to be investigated. Then URANS and LES fluid dynamics computations will be adopted to evaluate the performance of low-pressure turbine blades. In this phase the main part of data will be generated, and all these data will be the input for HPDA and AI applications.

Several data format will be adopted along the workflow coherently with the requirements of the different tools. In particular the following format will be adopted:

- ASCII non formatted file (overall input/output, AI database, etc.)
- binary files (STARCCM+ solution file, Ensight gold file format for Paraview post processing)
- HDF5 format may be also adopted for HPDA

As reported in the deliverable D1.3 Data Management Plan, the data collected during the ACROSS project, will be discriminated as confidential, hence not shareable, and others that will be public with the objective to promote understanding and comprehension about which type of knowledge can be extracted thanks to the advanced Computational Fluid Dynamics Analysis: a great insight to explore complex flow fields aimed at minimizing loss sources and get more efficient fluid processes.

3.2.9 Co-design Requirements

Concerning AI tools, in the proposed development plan, the WP5 team will directly work on the new tools, focusing on a more effective framework/languages TensorFlow/Python with respect to the current one (Fortran based) and it will be directly involved in the porting to GPUs, but some support from the WP3 will be necessary. The WP5 team will also support testing on ultimate generation accelerators (FPGAs, neuromorphic processor, etc.), by making available tools and databases for porting/testing.

Concerning HPDA/LES improvements, the new workflow will place new challenges for orchestration and data management. Obtaining a large speed-up, indeed, will require a proper HW and orchestration between the different tools. Even if it is likely that, in the improved version, LES calculation will be performed for a reduced number of iterations, to reach relevant benefit, it is important that the co-processing will have a negligible impact on the calculation speed. HPDA/LES improvements will be mainly in charge of the WP5 team that will develop the new procedure in which LES and HPDA interact while LES is running to minimize computational costs. At the same time, a contribution of the WP4 will be crucial focusing on the overall process orchestration in an advanced HPC infrastructure.
In the new workflow the cluster architecture and the way resources will be management will play a key role to ensure fast read/write operations and data exchange between LES and HPDA, and the correct execution of concurrent tasks. On this topic, indeed a large amount of data is written at run-time and must be effectively exchanged.
As a further challenge, it is important to remind that LES and HPDA may be carried out on different HW (CPUs for LES, CPUs/GPUs for HPDA).

Requirements summary:

- Baseline
  - URANS
    - One calculation
      - 10 hours
      - 10 cores
    - ~1000 calculations performed simultaneously
  - LES calculation
    - 500 cores
    - 10 days
  - HPDA
    - 500 cores
    - 10 hours
    - 1.5 TB
- Improvements
  - Introduction of the new DS into the industrial design procedure
  - Adoption and improvement of advanced tools in all the phases of the new DS
    - AI of the DS itself from hardware and software perspective
    - Integration of LES an HPDA from the orchestration perspective
- HW requirements
  - CPU compute nodes
  - Accelerators (GPU, FPGA, VPU, NNP, etc.)
- SW requirements
  - Commercial STAR-CCM+
  - Matlab, Fortran, Python, Java (HPDA)
  - TensorFlow/Keras
  - OpenVINO
  - Paraview
- Other requirements
  - AI related development of new tools
  - Porting to GPUs or other accelerators
  - Orchestration between different tools
  - Fast read/write operations and data exchange between LES and HPDA
  - Checkpointing

4 WP6 - Weather, Climate, Hydrological and Farming Pilot Requirements

4.1 Expected KPIs

WP6 aims to develop and demonstrate complex hydro-meteorological and climatological workflows and reach a high TRL to enable operational exploitation of the innovation developed in the context of the project.
In particular, the following KPIs, defined for this pilot, refers to IFS global-scale Numerical Weather Prediction model and WFLOW hydrological model.

KPI-2.1	Demonstrate Numerical Weather Predictions (NWP) workflow with IFS model resolution improved from the current 9.0km operational resolution to 5.0km.
KPI-2.2	Improve WFLOW (hydrological application) runtime performance when compared to today's capabilities at least factor 5, to enable full ensemble simulation
KPI-2.3	Demonstrate hydrological simulations over Rhine and Meuse basins (220.000 km2 area) adopting 1.0km model resolution

Table 3 KPIs for WP6

4.2 Current Performance

ECMWF is developing and executing operationally a global-scale NWP called Integrated Forecasting System (IFS). IFS is currently executed on dedicated hardware (ECMWF supercomputer - Cray XC40, Xeon E5-2695v4 18C 2.1GHz, Aries interconnect) twice a day in 1h critical window. During that timeframe, we execute one high-resolution (9km) simulation and 51 ensemble members at 16km resolution. The scalability of IFS and ICON models has been demonstrated prior to this project on multi-petaflop systems (e.g. Summit, JUWELS).

WFLOW baseline implementation in Python requires 20' for each model member, for a total running time of about 17 hours. The baseline version is a sequential code, thus not able to exploit any parallel computing resource.

WRF is a community model with a large user-base of more than 36000 registered users and a stable codebase, with compile options for all major CPU architectures. First versions optimized for GPU exploitation are appearing (i.e. AceCast from TempoQuest - based on proprietary code).
A tipical WRF domain configuration with the outer domain covering Europe, 2 nested sub-domains, with innermost domain at 2km and roughly 800x800 cells requires up to 1000 cores to generate a 48h regional downscale in 2h wall-clock time on IT4I Barbora supercomputer.

4.3 Expected Improvements

We will work on efficient exploitation of high performance data store based on non-volatile or flash memory, adding support to our FDB object store. In a time-critical complex workflow we experienced that one of the main issues is the contention among data writer (NWP and Climatological model) and readers (post-processing tools).

We aim at extending the ICON model, adding support for FDB object store, in order to improve model IO and to support semantic queries over the generated climatological datasets. A single high-resolution climatological simulation can produce output data exceeding a Petabyte. Data are produced as time-steps, and each time-step includes all the multi-dimensional meteorological fields. By contrast, most climatological analysis and data processing require time-series over a meteorological variable. In the context of the ACROSS project we will investigate the benefit of high-performance hierarchical data stores coupled with FDB object-store for the analysis of large-scale climatological simulations.

We aim at optimizing WFLOW implementation by porting the Python code to Julia, parallelising the existing algorithm and exploiting horizontal scalability of Cloud resources. The goal is to speed-up WFLOW execution forced by a full set of IFS ensemble members, and reduces wall-clock time to less than 1h in order to enable operational exploitation of WFLOW over Rhein and Meuse basins by the Dutch water-management authority. In the context of ACROSS we will also assess the benefits of direct connection of WFLOW to FDB object store co-located in the same datacenter, over the web-based interfaces for data fetching.

We expect an improvement of WRF and WRF-DA on the upcoming ACROSS computing resources (namely CINECA Galileo100 and IT4I Karolina) thanks to the more recent architecture of such new machines, and thanks to the exploitation of fast I/O devices. Moreover, we expect an overall improvement of the overall execution time in the workflows involving global scale NWP and regional downscaling, thanks to the exploitation of the ACROSS orchestration services.

4.4 Expected Performance

Current ECMWF operational weather forecasts are computed at 18km resolution (for the 51 ensemble members) and 9km resolution for the high-resolution deterministic run. In the context of ACROSS we aim at demonstrating IFS model execution at 5km or finer resolution. Depending on the computational resources made available through the ACROSS project, we will consider the deterministic HRES model and optionally a selection of the ensemble members. Moreover, we will demonstrate the full post-processing chain, responsible for user-defined product generation.

We aim at optimizing WFLOW implementation to reduce the overall execution time to <1h (compared to actual 17h), and thus enable operational exploitation of such model. We target ACROSS HPC and Cloud resources to demonstrate WFLOW achieved performances. as highlighted in the previous section, Deltares is willing to demonstrate operational readiness of WFLOW over large Basins such as Rhine and Meuse, and thus involve Dutch water-management authorities for further exploitation of the innovations introduced by ACROSS.

In order to provide accurate and timely farming advisory services, we are expecting to demonstrate NWP regional downscaling over Europe and Greece performed by WRF, enriched by data assimilation of Neuropublic weather stations. Exploiting the data interoperability provided by FDB, ACROSS infrastructure and the EuroHPC computational resources we are expecting to deliver ensemble regional downscaling with data assimilation in less than 2h wall-clock time.

4.5 Improved Workflow Description

In this section, we present the three workflows we aim at implementing in the context of ACROSS. Each workflow is represented as a simple chain of sub-workflows, to clarify model interactions, then each sub-workflow is expanded to better identify the computational tasks and the related dataflow.

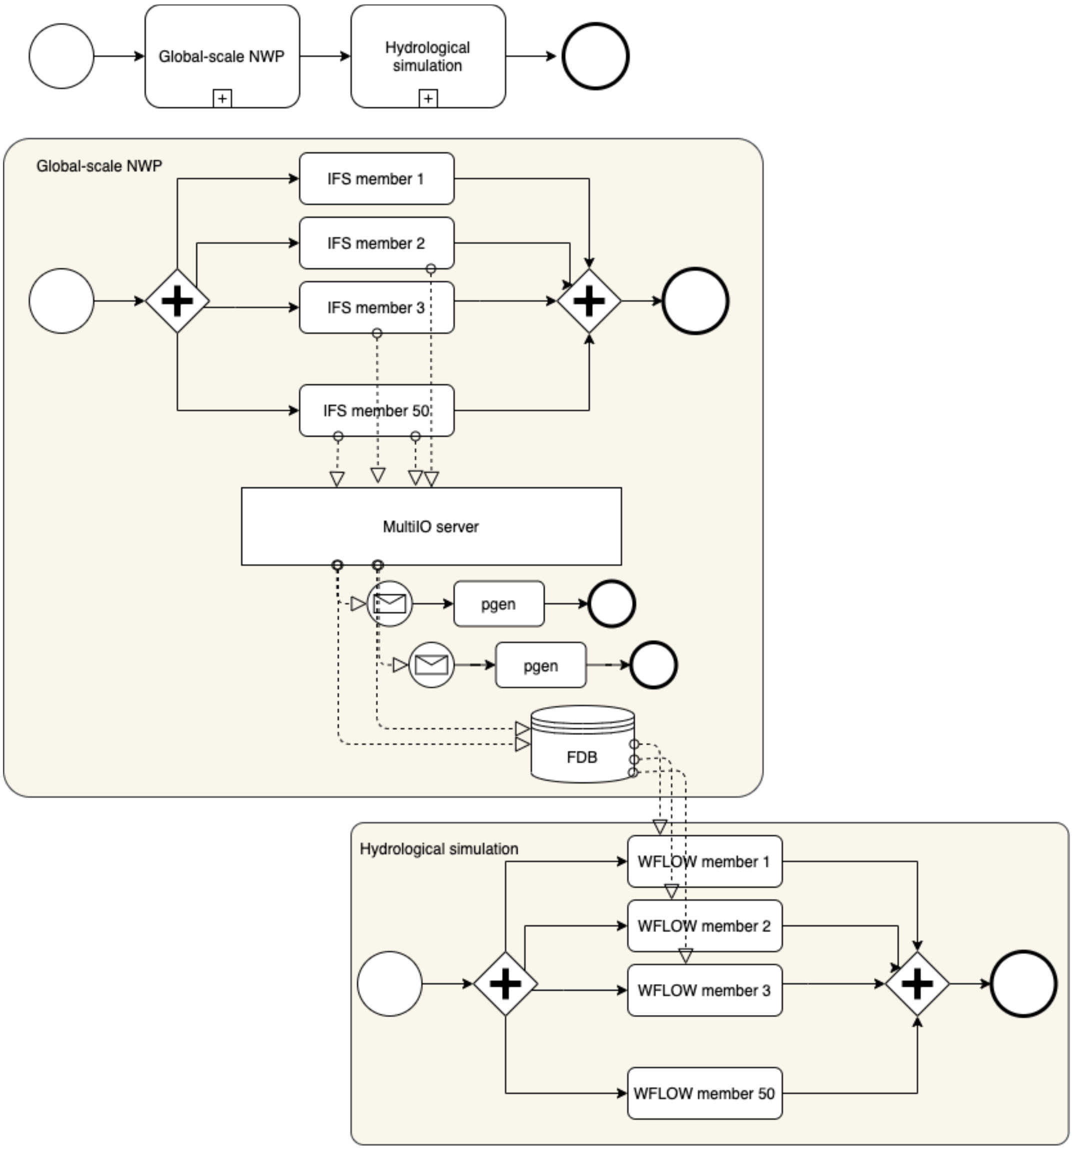

Figure 19 BPMN representation of Hydro-meteorological workflow

The Hydro-meteorological workflow is composed of 2 main steps: a global-scale Numerical Weather Prediction (NWP) workflow followed by a Hydrological simulation over Rhine and Muse basins.

The global-scale NWP and product-generation can be considered as a producer-consumers sub-workflow in which the IFS model is producing the global-scale forecasts and pgen is post-processing the global forecasts performing crop, re-grid and other user-defined processing to minimize data movement.
The sub-workflow mimics an operational setup that is usually executed 4 times a day in 1h time-critical window, but in the context of ACROSS is not executed in real-time or with a fixed periodicity.

The global-scale NWP requires two kinds of input data: static data (i.e. geographical datasets) and dynamic data (i.e. observational data). Both will be provided by ECMWF for selected test cases and are required to be available before the start of the workflow.

Generated data products are made available to Deltares WFLOW for the hydrological simulation. We foresee adopting both products generated by the HRES deterministic model run and the full set of ensemble members.

The hydrological model WFLOW uses static data (e.g., land use, soil, elevation information) which is generated using the Deltares HydroMT framework (available on GitHub). For dynamic data, the model is forced with precipitation, temperature and potential evaportranspiration. Dynamic data will be provided by ECMWF (for weather) and MPI (for climate).

ECMWF will run the IFS global NWP model on EuroHPC computing resources and output data will be made available to downstream applications through the FDB object-store. The different ensemble members of precipitation, temperature and potential evapotranspiration of the forecasts will be extracted for the Rhine and Meuse basin (north-western Europe). Each member will be used to force an instance of the wflow-sbm model.

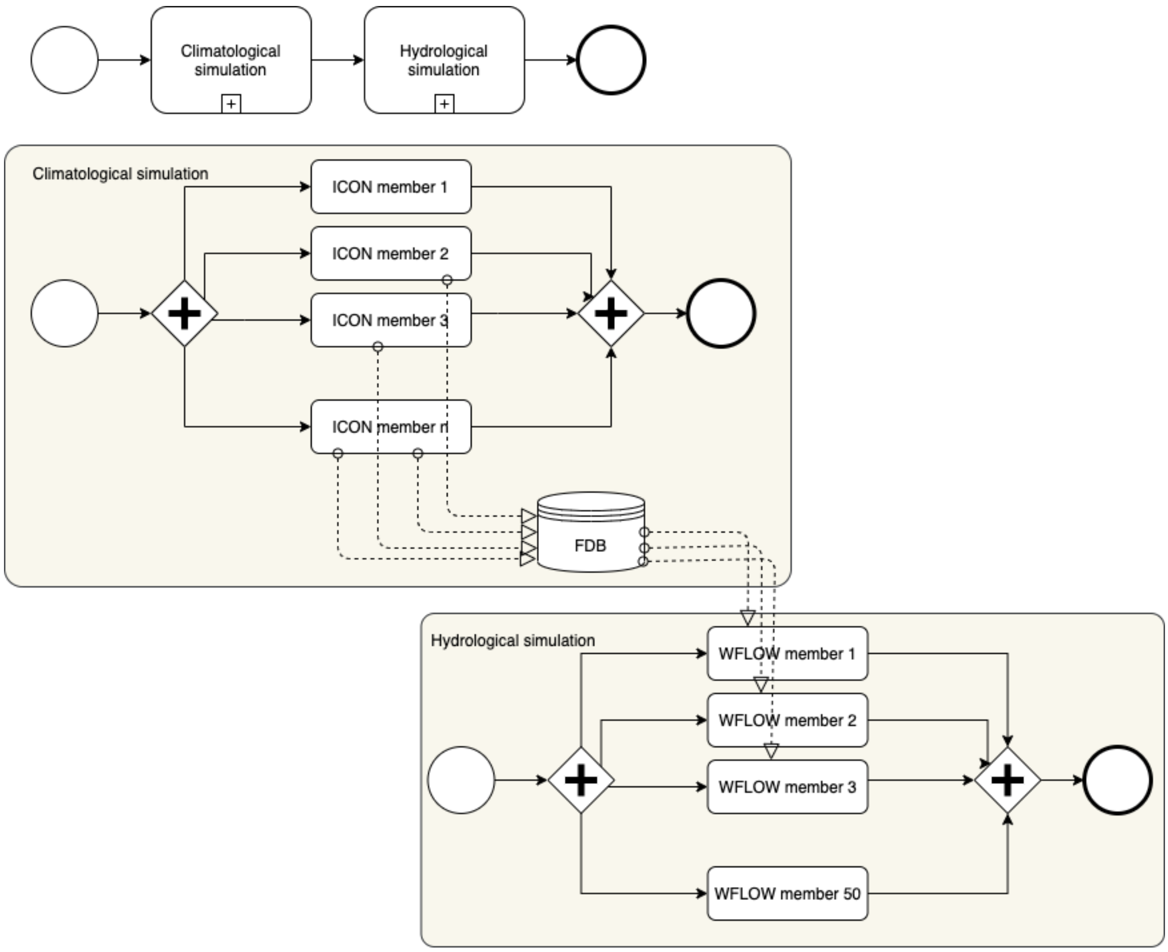

Figure 20 BPMN representation of Hydro-climatological workflow

A similar workflow will also be used for the climate model simulations, where MPI will generate various ensembles of the ICON model and make these available in the FDB. Each member will subsequently be used to force wflow-sbm on the Rhine and Meuse basin.

The hydrological part of the workflow is analogous with the previous workflow. However, when taking into account that the climate simulations consider time varying land-use, potential physical interfacing considerations need to be discussed. For simulations with respect to future climate scenarios the respective on(c)e have to be selected.

BPMN representation of Farming advisory workflow:

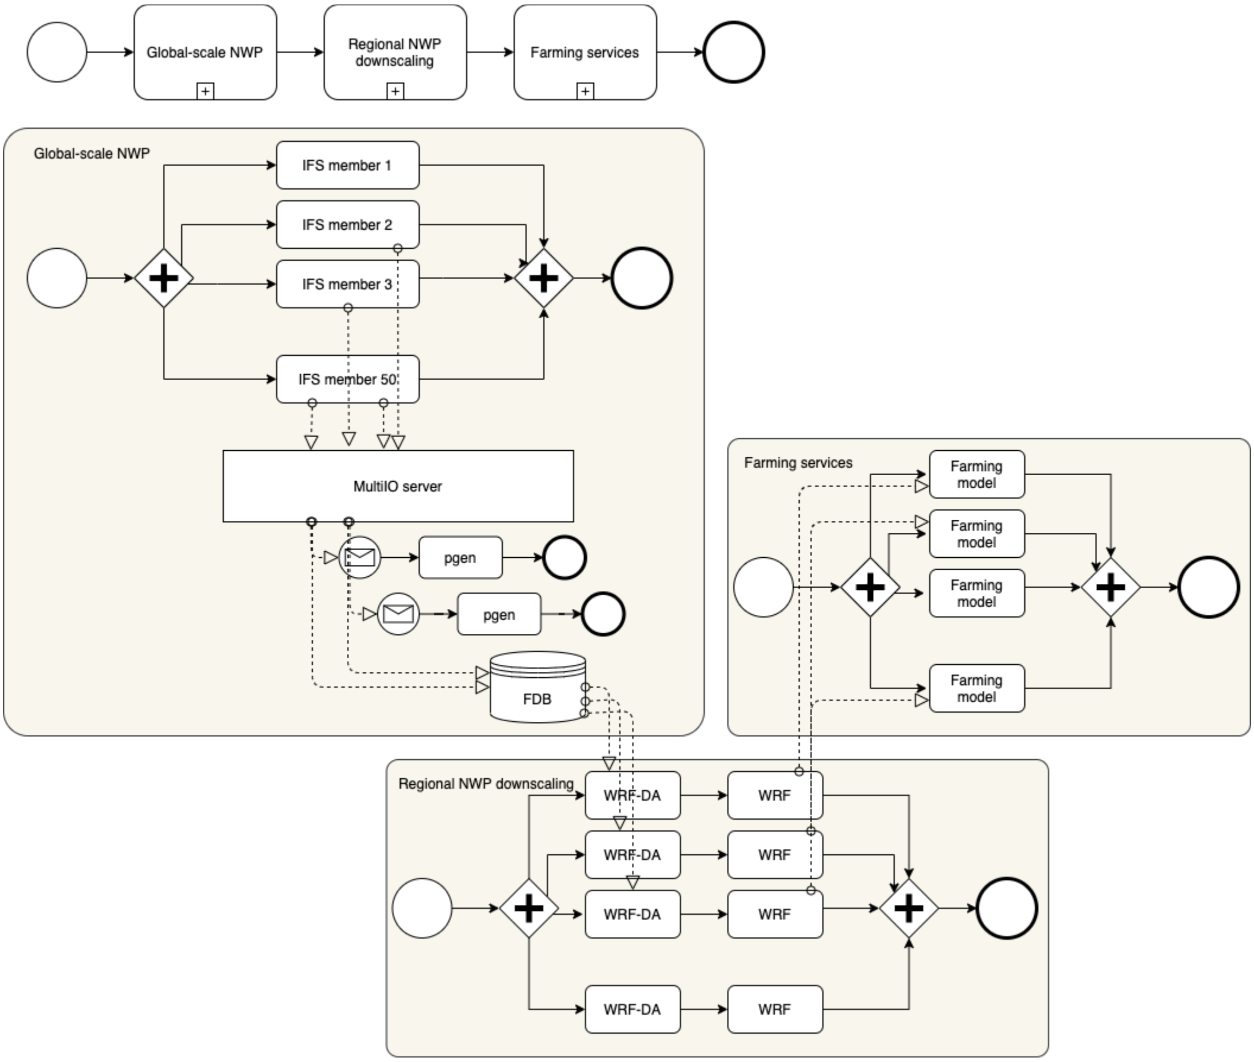

Figure 21 BPMN representation of Farming advisory workflow

The Farming advisory workflow is composed by 3 main steps: global-scale NWP and product generation, regional downscaling of NWP with data assimilation, and farming specific post-processing applications.

The global-scale NWP and product generation is analogous to the case of the first workflow.
We may consider to enrich the workflow with ML-based data analysis for identifying specific phenomena in the computed forecasts.

The ECMWF global-scale forecasts (either HRES and ensemble) will be downscaled for the regions of Europe and Greece using the WRF model. The downscaling procedure will also be enriched by assimilating NP's weather stations. Sensitivity tests will be conducted, in order to decide which schemes of WRF, concerning radiation, cloud microphysics, boundary layer etc. are suitable for the model's configuration over the area of eastern Mediterranean – Greece. The improved forecasting skill of the model can also be demonstrated by showing its benefits on smart farming necessities concerning the weather's impact on agriculture.

WRF forecast products by ACROSS infrastructure will serve as input for smart farming applications developed by NEUROPUBLIC (The smart farming applications will not be developed on the ACROSS infrastructure). Some examples (not mandatory though as it will be decided later) could be SMS texting concerning severe weather alerts, parts of mobile/pc applications (e.g. giving also results/forecasts of leaf wetness or evapotranspiration, etc).

4.6 Hardware Requirements

IFS Global-scale NWP:
Computing resources required by current implementation of IFS, product generation and FDB are CPU only. IFS model requires fast fabric (IB)
FDB has support for different storage technologies (Lustre parallel file-system, Ceph and PMEM). Lustre is used in operation, while the other two are experimental.
Each full model run (HRES and ensemble) generates ~70TB at operational resolution. We expect to reach 250-280 TB/run at doubled resolution.

ICON Climatological simulations:
Computing resources required by ICON are either CPU only or CPU/GPU setups. A fast fabric is required. Data amount is expected to be a minimum of 70 TB per simulated month dependent on the requirements of data users, thus a full-resolution climatological simulation over 30 years may require ~25PB. Depending on the time span expected for the climatological simulations, a significantly reduced data amount could be provided.
For FDB the overhead of the index is negligible compared to the size of the data.

Hydrological simulations :
Computing resources required by WFLOW are CPU only. Generally about 8-10 GB of RAM is needed to run the model at the scale of the river Rhine and Meuse. For the real-time forecast run, where the model is forced with IFS data, it is expected that about max 1-2 GB of data/run are simulated depending on whether gridded information will be stored. In case only the time-series at specific location are stored only up to 10-20MB per run is generated.
For the long-term climate simulation, in the order of 1-3 TB of data are expected to be generated depending on which gridded product will be stored.

WRF regional-scale NWP downscaling:
Computing resources are CPU. For deterministic and ensemble forecasts (taking into account also the different configurations of the model, number of case studies, ensemble members, etc.) 50 TB of storage will be needed (it is an estimation though, the number may change).

4.7 Software Requirements

IFS, ICON and WRF are usually compiled for Linux and have rather similar requirements.
The overall list of dependencies includes:
bash, perl, make, m4, Python, BLAS/LAPACK, netCDF libraries (netCDF, HDF5, pnetcdf), GRIB libraries (libaec, eccodes, JasPer, libpng, zlib).

WFLOW uses the Julia programming environment, which can be operated standalone or be converted into a Docker image to enable operation in the Cloud. In this last application, it will need both Docker, Kubernetes and the ARGO workflow manager.

Fortran 2008/C/C++11 compilers (both GNU and Intel are fully tested) with OpenMP and MPI.

4.8 Data Management

Global-scale and climatological simulations at high resolution are the most demanding applications for data management. We foresee to produce 250-300 TB in 1h of 5km ensemble global-scale NWP. We are working on GRIB compression based on libaec to reduce such requirement.
NWP output post-processing (product generation and regional downscaling) is going to access nearly 70% of the generated output data within the time-critical execution time-frame.
Grand-ensemble climatological simulations are less intensive but are going to produce huge output dataset in the order of several PiB.

4.9 Co-design Requirements

Due to the size of the data produced by the WP6 workflows, we foresee the need to carefully tune the configuration of multi-layer data stores available on ACROSS computing resources, to efficiently support data management requirements. Moreover, the execution of tightly coupled workflows, with immediate exploitation of global-scale NWP outputs by subsequent models in the processing chain poses a challenge in terms of data contentions.

Requirements summary:

- Baseline
  - Global-scale NWP called Integrated Forecasting System
    - Running on dedicated ECMWF supercomputer - Cray XC40, Xeon E5-2695v4 18C 2.1GHz, Aries interconnect
  - WFLOW
    - single core, 17h walltime
  - WRF
    - 1000 cores, 2h walltime
- Improvements
  - efficient exploitation of high-performance data store based on non-volatile or flash memory
  - extending the ICON model with support for FDB object store
  - optimizing WFLOW implementation by porting the Python code to Julia
  - improvement of WRF and WRF-DA
- HW requirements
  - CPU compute nodes (8-10 GB RAM)
  - GPU accelerator
  - IB interconnection
  - Lustre parallel file-system/Ceph/PMEM
  - From 70 TB to 250-280 TB/run (meteorological runs)
    From 70 TB to multi-PB (climatological runs)
- SW requirements
  - WFLOW
  - Dependencies (bash, perl, make, m4, Python, BLAS/LAPACK, netCDF libraries, GRIB libraries)
  - Fortran 2008/C/C++11 compilers
- Other requirements
  - configuration of multi-layer data stores
  - processing chain’s data contentions solutions

5 WP7 - Energy and Carbon Sequestration Pilot Requirements

5.1 Expected KPIs

The following KPIs are defined for this pilot:

KPI-3.1	Improving OPM Flow runtime performance scaling when compared to today’s parallel capabilities, scaling to 1000 processes with reasonable efficiency.
KPI-3.2	Carrying out flow simulations on large grids for long-term migration scenarios (> 1000 years), on models with up to 100M cells.
KPI-3.3	Running direct flow simulation on models consisting solely of processed seismic data, at high resolution, with automatic and dynamic coarsening/refinement, on models with up to 100M cells.
KPI-3.4	Demonstrating analysis of simulation results in-situ using methods from the AI spectrum in 3 new workflows. Increase by 50% the overall data processing throughput (i.e., the number of scenarios

	evaluated per unit of time and the requests per second served in extreme cases).

Table 4 KPIs for WP7

5.2 Current Performance

We have successfully run OPM Flow on clusters up to 256 cores in parallel (multi-node). We observed good scalability up to around 100 cores, and poorer scalability after that. 256 cores were no faster than 128. These results vary with case and hardware. Scalability is usually poorer on non-HPC systems such as workstations or laptops, as can be expected. The amount of interaction between wells also has strong impact on scalability beyond 100 cores, since long well paths that intersect large parts of the domain reduce the quality of partitioning that is possible to achieve and increase connectivity between the regions of the domain. We expect CO2 storage cases to be relatively benign in this respect though, as there is less well interaction. A plot of experimental scaling results is provided in Figure 22. A best-fit exponential (orange dashed line) gives a coefficient of almost -1, which is the ideal scalability (green line).
For the seismic cube use case, we have run trials with up to 6M cells. This represents a rather small case in this context and was done as a proof of concept.

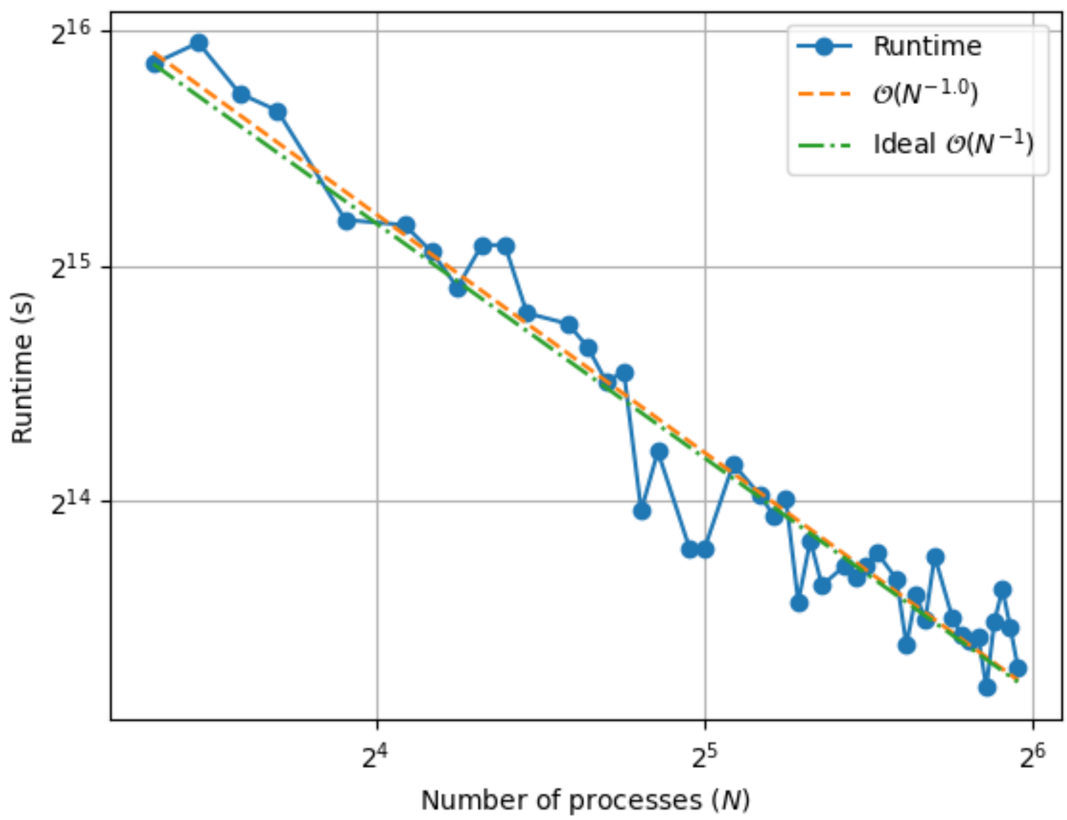

Figure 22 Scalability of a Sleipner case variant on the Saga HPC system

Damaris was experimented and validated up to 14,000 cores on top supercomputers including Titan, Jaguar and Kraken. It showed an increase of the sustained write throughput by a factor of up to 15 compared with standard I/O approaches; for example, it allows almost perfect scalability of the simulation up to over 9,000 cores, as opposed to state-of-the-art approaches that fail to scale. It enables a seamless connection to visualization software (Visit/Paraview) to perform in situ analysis and visualization in a way that impacts neither the performance of the simulation nor its variability.

5.3 Expected Improvements

OPM Flow will be improved in order to provide better parallel scaling, in particular in the area of parallel output. OPM's use of GPU accelerators will be improved. ERT will be improved to take advantage of advanced orchestration capabilities. Adaptive coarsening and refinement of computational grids will be implemented in OPM Flow, to enable direct simulation on seismic cube data with high efficiency.

5.4 Expected Performance

We expect to be able to scale single-case runs of very large cases (20M-100M cells) to run on up to 1000 cores with reasonable efficiency. We expect output to not be a bottleneck for such runs. We expect ensembles of up to 100 members to be run concurrently or semi-concurrently with high efficiency and utilization of allocated resources. For the seismic cube use case, we expect to encounter cases with up to 100M cells, that size therefore represents the upper case size limit that we consider in this project.

5.5 Improved Workflow Description

The overall workflow is described in Figure 23, with more detail in Figure 24 for the process of running a single simulation case. Figure 24 also takes into account the expected improvement from using Damaris to make the output parallel. The workflows are typically run in batch mode, with no real-time constraints. Several (3-7) iterations of the main history matching (HM) loop must be run sequentially: you cannot start a HM iteration until the previous one is complete (although it can be sufficiently complete even if not all cases of the ensemble have been successfully run). Note that each individual "Run single ensemble case" task can be an individual job, and that they are independent of each other. Typically, 50-100 such tasks will be run per HM iteration. One can consider progress by:

- Current iteration number in the HM loop.
- Within the current HM iteration, the number of single cases from the ensemble that have been run so far.
- Within a single case run, how many timesteps or how much simulation time has been completed (requires not treating the simulator as a completely black box).

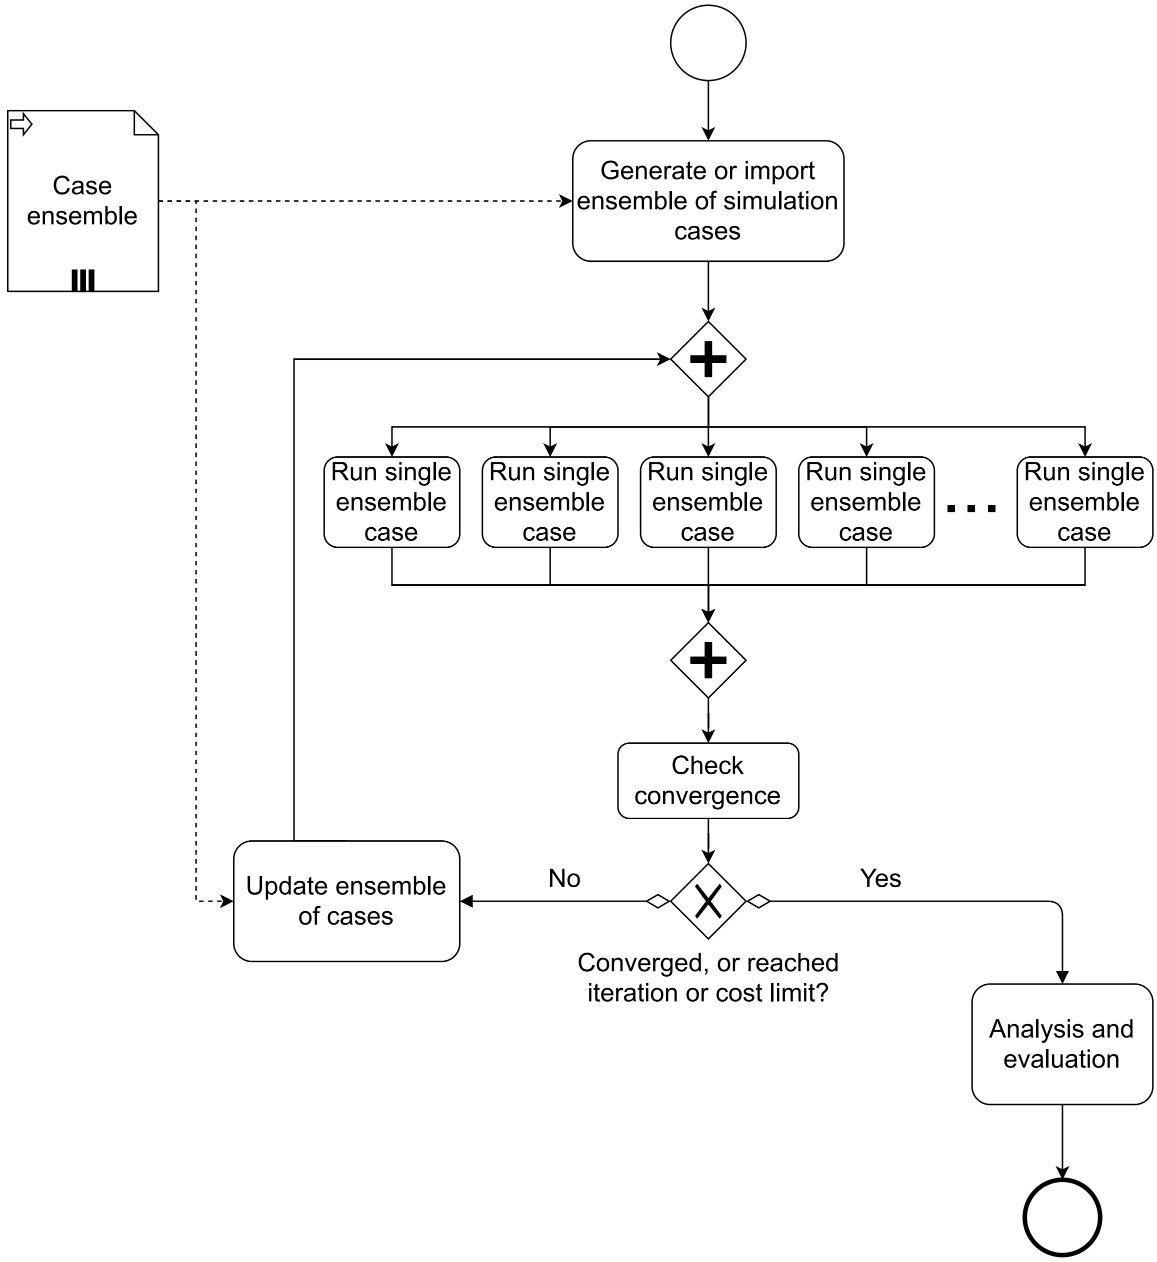

Figure 23 Workflow graph for history matching (HM) workflow Main HM loop tasks are the "Update ensemble of cases", multiple parallel "Run single ensemble case" and "Check convergence" tasks.

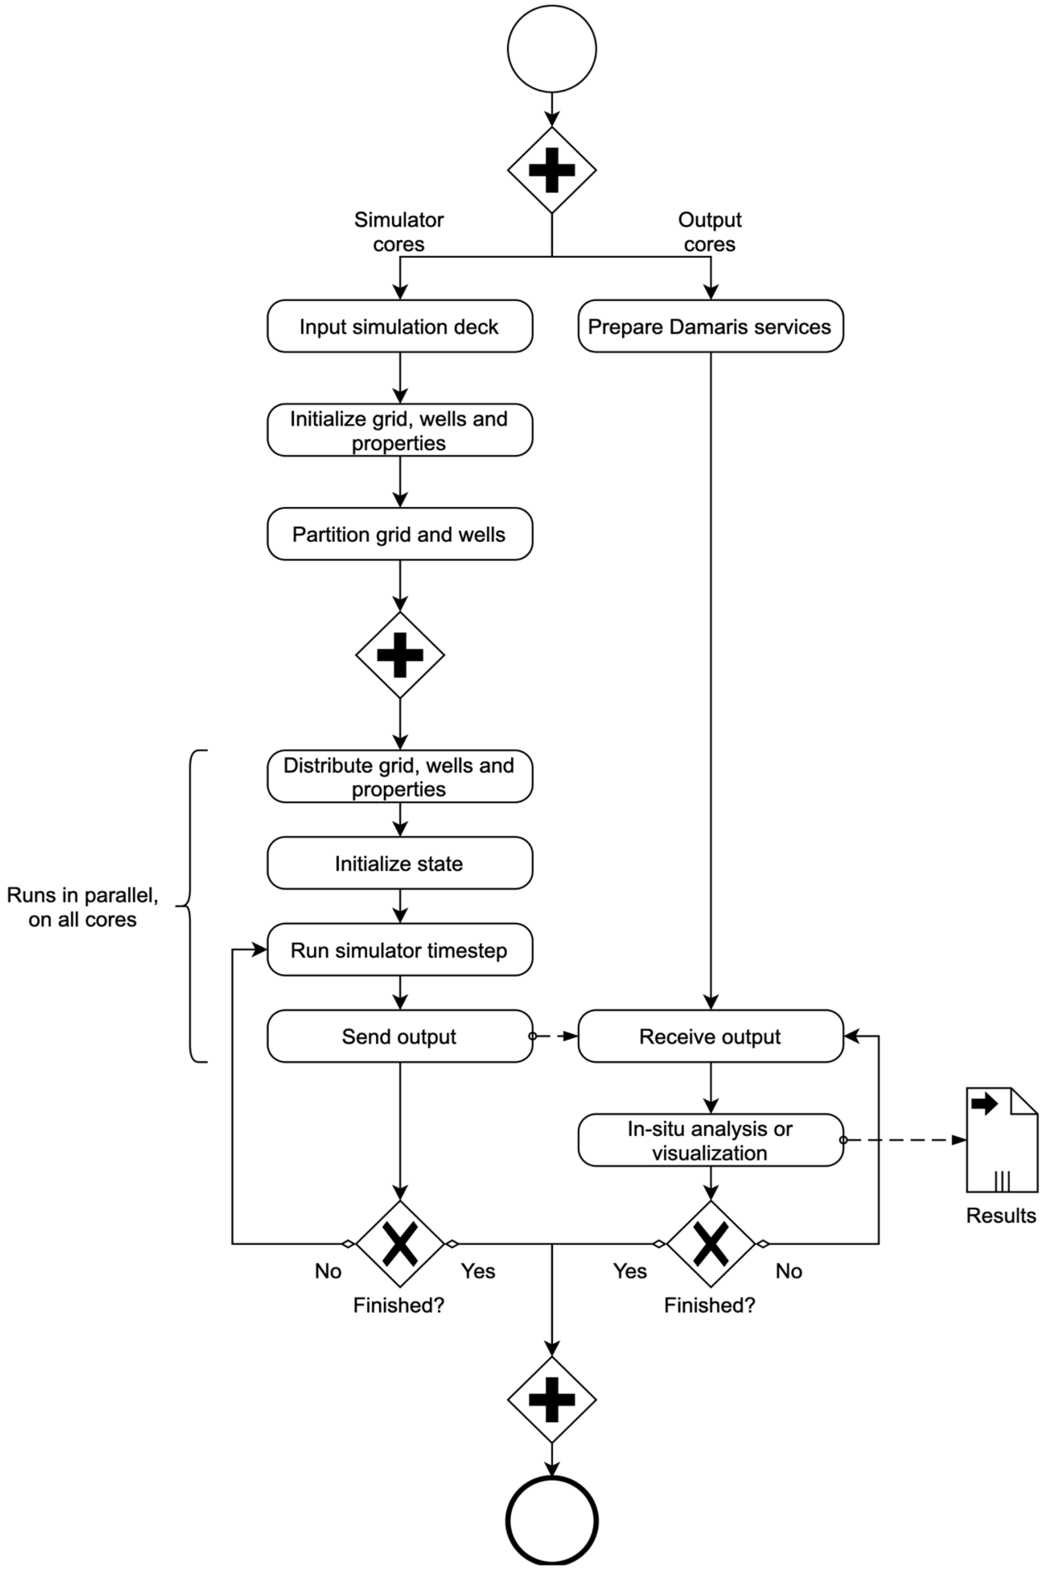

Figure 24 Workflow graph for the "Run single ensemble case" part of the HM workflow shown in Figure 23. Note that at the moment, memory usage is very high on a single simulation core during the initialization process, as that one core reads and processes all the input data.

5.6 Hardware Requirements

Processors: At this point both the simulator OPM Flow and the ERT tool need to run on CPUs, although parts of the simulator can be run on GPUs experimentally, with a restricted set of linear solver options.
CPU-time and memory: running an example case with 17M cells requires approximately 35 minutes on 64 cores, with two threads on each core. Memory usage for each core is approximately 2GB, in addition, one of the cores must have significantly more memory, approximately 40 GB. The reason for this is that one single process handles all input and initialization at the start of the run, which consumes significant memory.
Investigation of optimal computational load per node and related memory sizes in conjunction with the schedule/placement procedures is part of our approach in improving the performance.

5.7 Software Requirements

The following software is required for this use case:

- OPM Flow
  - In active use for industrial exploitation
  - Open source (GPLv3+)
  - Parallel paradigms are MPI for domain decomposition, with OpenMP used in parts (assembly of linear system)
  - GPU acceleration exists experimentally, but only for linear solver parts
  - Written in C++, compiled with gcc or clang (also small bits in C, some Python)
  - Developed on Linux (mostly) and macOS, can also be used easily with Windows Subsystem for Linux.
- ERT
  - In active use for industrial exploitation
  - Open source (GPLv3)
  - Can start jobs directly or using SLURM
  - Written in Python 3
  - Developed on Linux

The program OPM Flow reads input specifying a computation grid with associated petrophysical properties (permeability, porosity etc.), fluid properties (relative permeability, capillary pressure, viscosity, density, mixing properties etc.), wells and facilities, initial state of the reservoir, and a schedule for the operation, control or response of the wells. The output consists of 3D fields for the reservoir state (pressure, saturations, fluid composition) and time series (for example well injection or production rates).

OPM Flow simulates porous media flow by solving PDEs for the reservoir behaviour numerically, coupled to well or facility models. The PDEs are solved by Finite Volume methods in space, with an implicit Euler time discretization. The resulting nonlinear system of equations for each time step is solved by a Newton-like method. The most computationally intensive parts of OPM Flow are the linearization stage, where the nonlinear equation residual and Jacobian matrix is computed, and the linear solver stage. The Jacobian is computed by the automatic differentiation (AD) method, applied locally in a high-performing manner that keeps the flexibility of the AD approach. The linear systems are solved by default using a block-BiCGStab method where one block A_{ij} of the matrix contains the derivatives of the equations for cell i with respect to the unknowns of cell j. This is preconditioned either with block ILU0, or with a two-stage CPR preconditioner that first extracts and solves a pressure-type equation with AMG, then applies ILU0 to the full system.

The program ERT reads input specifying a history matching ensemble, and parameters for the history matching process. It performs the history matching process by repeatedly running the reservoir simulator (OPM Flow) as a black-box task, with varying input. The output consists of updated ensembles of cases (each such case is an input case for OPM Flow) and diagnostics information.

The future version of OPM Flow is intended to include the Damaris middleware to handle parallel output. Damaris provides scalable parallel output by reserving a few cores or nodes for collecting and creating output. The properties of Damaris are:

- Used for industrial exploitation.
- Open Source (LGPL v3)
- MPI used for parallelization (v2+ required)
- Written in standard C++11, compiled with gcc.

- Requires POSIX or SysV unix system libraries, XSD for XML input, Xerces-C and Boost libraries (log, date_time, system, filesystem)
- Optionally, VTK/ParaView or VisIt for visualization, HDF5 for output format.
- Developed on Linux.

For ease of deployment and updating of software it is very beneficial to have support for containerized applications, that can still exploit the hardware of the HPC systems, such as fast interconnects or GPUs. Singularity is at the moment the preferred container technology, although Docker is also useful.

For visualization purposes, accelerated VNC sessions that can access the output storage area of the simulations would be useful.

5.8 Data Characteristics

Data is unstructured, non-distributed, and for the current version written by single process. Data are usually confidential (both input and output). For testing, non-confidential data will be provided. Data sizes and corresponding estimated requirements are summarized tentatively in Table 5. Note that output sizes and computation times are estimated for a 20-year run, which is fine for history matching, but that much longer runs may be required, on at least a subset of each case ensemble, to study the long-term behaviour of the stored CO2.

Name	Description	#cells	Output size	CPU-hours
Norne	Oil and condensate field.	45k	188MB[1]	0.2
Sleipner CO2	Injecting 1Mton/yr for 20 years.	2M	0.5GB[2]	4
Refined Sleipner	Hi-resolution version of above.	18M	4.5GB[2]	40
Large Norwegian case	Industrial storage candidate site.	20M	5GB[2]	60
Huge multi-site model (concept, does not exist yet)	For example, the complete Utsira aquifer. May have Gton-level storage capacity.	100M or more	Unknown	Unknown

Table 5 Example data sizes

Storage is in regular local files currently. Need not be encrypted or compressed, but if the features are available, they may be useful. In the future, with in-situ visualization or analysis, data movement latencies may become important.

For the seismic cube use-case, data is not unstructured, but semi-structured voxel data (where each point in a regular 3D grid has one discrete data point) for the finest level of resolution, with up to 100M voxels. With the improved OPM Flow supporting adaptive computations, runtime and memory needs are expected to be lower than for regular cases of similar size run without exploiting such adaptivity.

5.9 Co-design Requirements

For ERT to be able to start the single-case runs and monitor their completion, an API for orchestration must be adopted, improved, or created. This can then be used from ERT as an alternative to directly starting jobs from ERT with (for example) SLURM, thereby enabling an orchestration component to optimize this part of the workflow.

Requirements summary:

- Baseline
    - OPM Flow
        - Tested up to 256 cores
        - Good scalability up to 100 cores
        - Example case with 17M cells

[1] *This is for compact domain-specific output (Eclipse binary format). VTK output is 4.1GB.*

[2] *Eclipse binary format, yearly output for 20 years.*

            - 35 minutes runtime
            - 64 cores
            - 2 GB RAM per core + one 40 GB RAM core
    - Seismic cube case
        - 6M cells proof-of-concept tests
    - Damaris
        - Tested up to 14.000 cores
        - Good scalability up to 9.000 cores
- Improvements
    - OPM Flow parallel scaling improvement
    - OPM Flow support for GPU accelerators
    - ERT to support advanced orchestration capabilities
    - Adaptive coarsening and refinement of computational grids implemented in OPM Flow
- HW requirements
    - CPU compute nodes
    - GPU accelerators
- SW requirements
    - OPM Flow GPLv3+ (C, Python)
    - ERT GPLv3+ (Python3)
    - Damaris LGPLv3 (C++11)
    - VTK/ParaView or VisIt
    - HDF5
    - Singularity
    - VNC
- Other requirements
    - Adopt/prove/create an API for orchestration

6 Conclusions

The D2.1 aims to clarify the co-design pilot requirements as they emerged from the questionnaires filled by each pilot partners, WP5, WP6 and WP7. Specifically, the document describes the current status and the expected improvements for each use-case. As we have already reported in WPs (5-7) Pilot Requirement sections, these are the most relevant co-design requirement for each WPs:

WP5 - Greener aero-engine modules optimization Pilot Requirements: The main objectives set by GE Avio Aero are to reduce the productivity target (reduction of time-to-design with respect to the current situation) by at least 50%, acting on both the modelling aspects of the physical problem and the optimization of the hardware. The other aspect is to improve the quality of the numerical results, getting as close as possible to the experimental reference results. Improvements of several aspects of the overall procedure (elimination of I/O issues and waiting times) will result in improved workflows. Moreover, the target of this activity is to develop an innovative, data-driven, AI-powered DS for turbines, capable of switching from standard aero components to innovative, additive-enabled aero components, aimed at improving the efficiency of the low-pressure engine modules.

WP6 - Weather, Climate, Hydrological and Farming Pilot Requirements: demonstrate Numerical Weather Predictions workflow with IFS model resolution improved from the current 9.0km operational resolution to 5.0km. Improve WFLOW (hydrological application) runtime performance when compared to today's capabilities at least factor 5, to enable full ensemble simulations. Demonstrate hydrological simulations over Rhine and Meuse basins (220.000 km2 area) adopting 1.0km model resolution. Carefully tune the configuration of multi-layer data stores available on ACROSS computing resources, to efficiently support data management requirements.

WP7 - Energy and Carbon Sequestration Pilot Requirements: Improve OPM Flow runtime performance scaling when compared to today's parallel capabilities, scaling to 1000 processes with reasonable efficiency. Carrying out flow simulations on large grids for long-term migration scenarios (> 1000 years), on models with up to 100M cells. Running direct flow simulation on models consisting solely of processed seismic data, at high resolution, with automatic and dynamic coarsening/refinement, on models with up to 100M cells. Demonstrate analysis of simulation results in-situ using methods from the AI spectrum in 3 new workflows. Increase by 50% the overall data processing throughput (i.e., the number of scenarios evaluated per unit of time and the requests per second served in extreme cases). All the pilots have presented a well-planned co-design requirements and improvements in order to define new strategies aimed to improve the efficiency of the complex workflow execution. This step is essential for the successful execution of the next deliverables, specifically, D2.2, devoted to the "Description of key technologies and platform design", and the achievement of the Milestone 1 "Awareness of project objectives and requirements".

References

[1] S. S. Haykin, *Neural networks: a comprehensive foundation*, 2nd ed. Upper Saddle River, N.J: Prentice Hall, 1999.

[2] D. Bertini, L. Mazzei, A. Andreini, and B. Facchini, "Multiphysics Numerical Investigation of an Aeronautical Lean Burn Combustor," in *Volume 5B: Heat Transfer*, Phoenix, Arizona, USA, Jun. 2019, p. V05BT17A004. doi: 10.1115/GT2019-91437.

[3] S. Paccati, D. Bertini, L. Mazzei, S. Puggelli, and A. Andreini, "Large-Eddy Simulation of a Model Aero-Engine Sooting Flame With a Multiphysics Approach," *Flow Turbulence Combust*, vol. 106, no. 4, pp. 1329–1354, Apr. 2021, doi: 10.1007/s10494-020-00202-5.

[4] F. Rubechini, M. Giovannini, A. Arnone, D. Simoni, and F. Bertini, "Reducing Secondary Flow Losses in Low-Pressure Turbines With Blade Fences: Part I — Design in an Engine-Like Environment," in *Volume 2B: Turbomachinery*, Phoenix, Arizona, USA, Jun. 2019, p. V02BT40A020. doi: 10.1115/GT2019-91280.

[5] M. Giovannini, F. Rubechini, M. Marconcini, A. Arnone, and F. Bertini, "Reducing Secondary Flow Losses in Low-Pressure Turbines: The 'Snaked' Blade," *IJTPP*, vol. 4, no. 3, p. 28, Aug. 2019, doi: 10.3390/ijtpp4030028.

[6] P. J. Przytarski, A.P.S. Wheler, D. Lengani, and D. Simoni, "Data-driven analysis of high fidelity simulation of multi-stage compressor," presented at the European Conference on Turbomachinery Fluid dynamics & Thermodynamics, Gdansk, Poland, Apr. 2021.