NEWSLETTER #8

FEBRUARY 2024

ACROSS
HPC Big DAta ArtifiCial intelligence cross Stack PlatfoRm TOwardS ExaScale

INSIDE THIS ISSUE

Welcome to the ACROSS Newsletter #8!

In the eighth edition of the ACROSS Newsletter we are going to present the most notable achievements. As such, the co-design approach is presented, followed by an overview on the main outcomes of technical Work Packages (namely WP3 and WP4).

ACROSS consortium has spent large effort in the validation process, which is an essential part for assessing the real capabilities of a designed system. In this sense, Pilots are the main providers for real applications that can be used for validating the whole platform.

Want to know more? Keep up to date on all project news @across_project and sign up for the Newsletter so you never miss out.

Questions? Contact info@acrossproject.eu

This 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.

Co-designing a cross-domain workflows' execution platform

The developed ACROSS platform is the outcome of co-design activities performed throughout the ACROSS project while integrating inputs and requirements from both technical and pilot WPs (see Figure 1). WP2 led the co-design activities of the platform's building blocks, WP3 focused on the infrastructure layer and utilization of acceleration technologies and specialized HWs, WP4 developed a custom smart orchestration solution, while the three pilot WPs, WP5, WP6, and WP7, provided the initial platform's requirements in order the pilots to be supported in their respective objectives.

A technical overview on the project

From the technical point of view, the ACROSS platform consists of the ACROSS SW stack and the underlying available computing and data infrastructures. ACROSS SW stack was created by interconnection of identified SW components and tools that were identified as the core SW building blocks during the initial requirement phase of the project. This ACROSS SW stack is portable and deployable on a number of different infrastructures, as proven by the current deployment scheme of the entire platform that consists of 2 production HPC infrastructures (IT4I's Karolina system, CINECA's G100 system), 2 customizable development environments, and 1 experimental HPC infrastructure (ATOS' NOVA system).

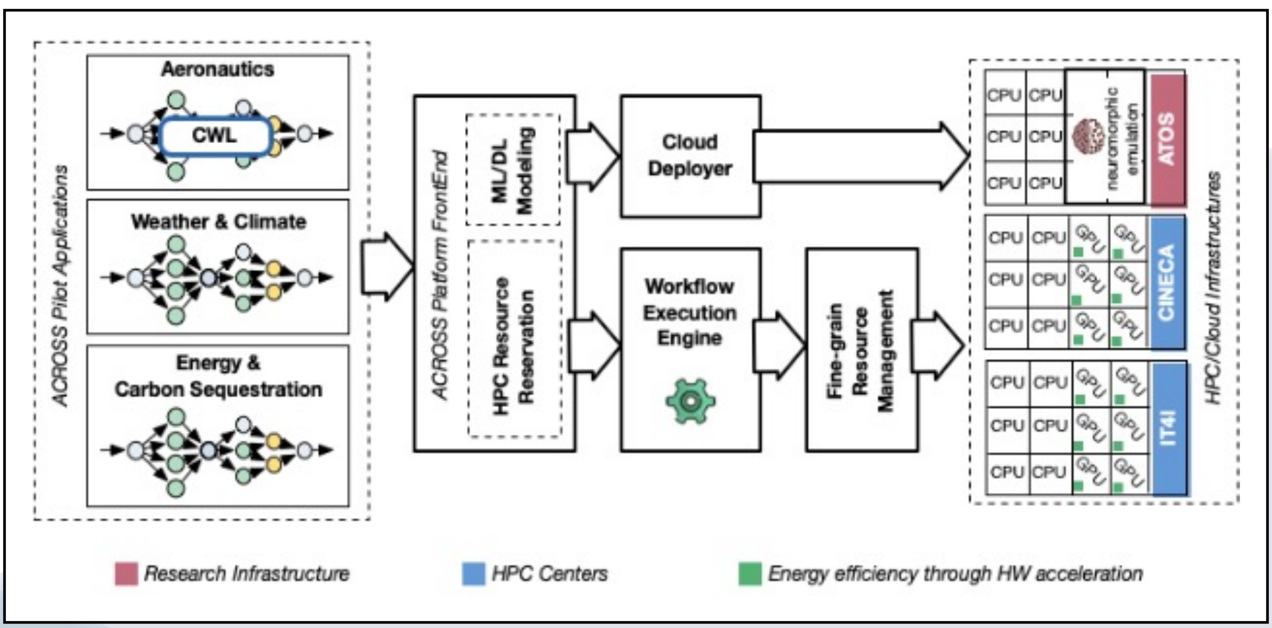

Fig. 1 – The ACROSS Platform – simplified view

The modular architecture of the platform supports the future development and extension of the platform's capabilities which opens several possibilities for the integration/interconnection of ACROSS platform with other platform/service providers. Based on the actual requirements, only a subset of the available ACROSS SW stack modules can be deployed to a specific HW location, thus enabling a wide range of configuration options.

HPC Technology Advancement

WP3 proceeded bottom-up, from the HPC hardware and software technologies to carry out technical developments that lay at a variety of abstraction layers, provided assistance for an optimal exploitation of advanced HPC technologies and investigated new avenues, especially in the field of AI.

Therefore, WP3 worked in close collaboration with other application work-packages, with a focus on the combined use of programming models (e.g., Nvidia CUDA-aware MPI for multi-GPUs exploitation), the parallelization and acceleration on GPU of Linear Algebra (LA) algorithms (e.g., solving large sparse linear systems using a variety of preconditioning methods, singular value decomposition (SVD) of large matrices, etc.). It also focused on the optimization of Artificial Neural Network (ANN) models, as well as on their acceleration, respectively, in training by GPU and in inference by OpenVINO™.

Beyond these classical actions, WP3 also investigated the future of AI through Neuromorphic Computing for which purpose a technical and methodological path was established to connect classic ANNs to advanced Spiking Neural Networks (SNNs) simulators and emulators.

Notably, the ATOS experimental NOVA cluster is a concrete demonstration of some of these efforts (see Figure 2).

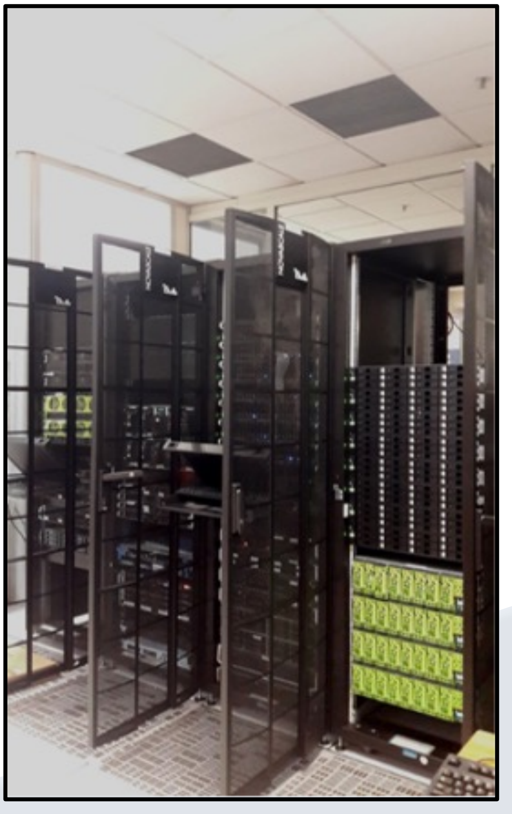

Fig. 2 – The experimental Atos NOVA cluster

Orchestration Software Stack

The WP4 focused on the design and implementation of the software stack that allows to effectively manage the execution of users' workflows on supercomputing resources made accessible within the project timeframe. In this context, the workflows focused on by the WP4 are the one structured with a mix of numerical simulations, high-performance data analytics and machine learning operations (considering both the training and inference phases); we refer to these as (complex) heterogeneous workflows. To manage their execution on modern supercomputers we architected a complex software stack (see Figure 3) which fulfills three main objectives:

- Providing a unified manner of describing the different steps composing the workflows and their targeted execution resources, even through an expressive DSL.
- Providing a set of advanced mechanisms for allocating compute resources.
- Exploiting better the underlying compute resources with a fine-grain workload allocation.

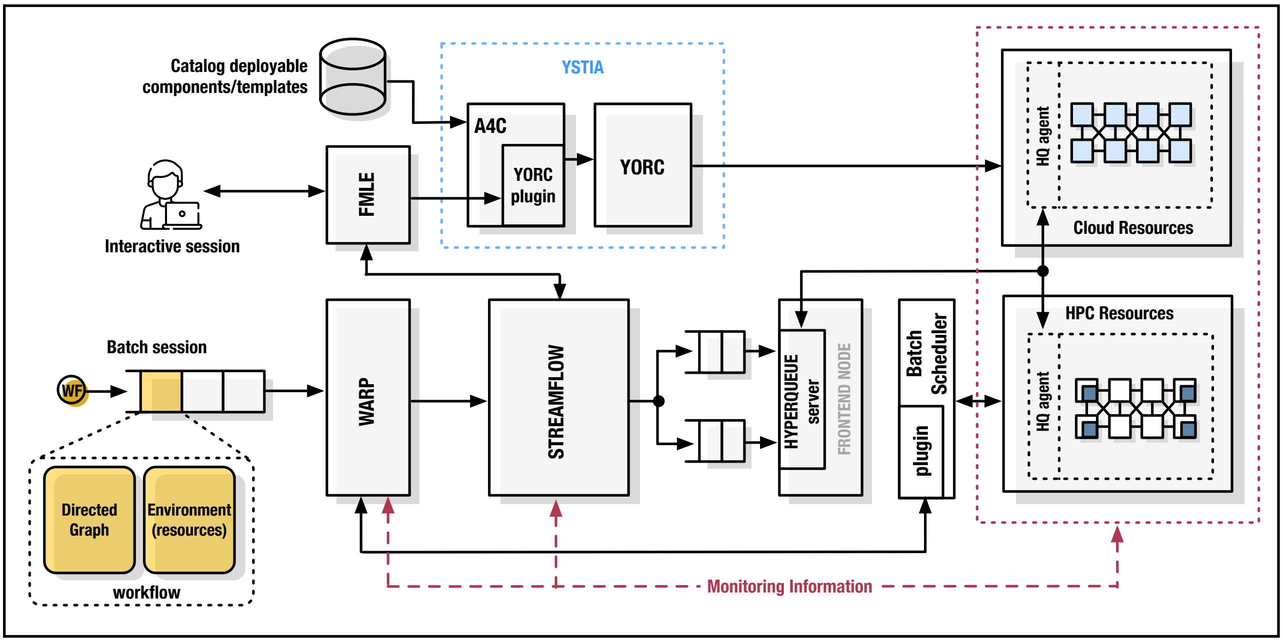

Fig. 3 – The ACROSS Orchestration Software Stack

The orchestration features

Specifically, WP4 introduced the WARP component to staging input workflows and getting the proper number of resources through a mechanism to create on-demand reservations. It also integrated a powerful workflow execution engine (StreamFlow) which comes with the CWL workflow description language (a standardized DSL for describing workflows). StreamFlow has been extended to include even the support for loops, as well as the CAPIO library that provides an extremely flexible way for supporting streaming applications. Then, the integration of the HyperQueue system, allowed the fine-grain exploitation of the (heterogeneous) compute resources (e.g., CPU cores and GPUs). Finally, the FMLE module along with the Ystia components provided the means for managing the machine learning training campaigns directly targeting the use of HPC nodes. We successfully experimented with this complex stack by running pilots' workflows.

The validation testbed

Aiming at demonstrating the capabilities of the toolchain composed by the integration of innovative software components, i.e., WARP, StreamFlow and HyperQueue, a testbed cluster has been used.

Being backed by high-performance bare metal compute node, the CINECA ADA Cloud based virtual cluster has been used to perform some experimental evaluation of the advantage offered by the on-demand reservation mechanism. The cluster (see Figure 4) was composed of 3 virtual machines running software services, and 14 workers, for a total of 144 vCPUs and 3.0TiB of shared storage.

The validation process considered the execution of a downscaled version of the Turbine use case where a LES simulation was followed by a HPDA analysis.

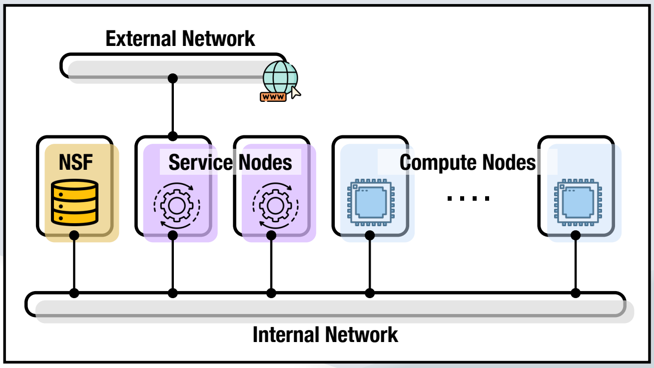

Fig. 4 – The testbed used to validate the software stack

Aeronautics Pilot

The aim of WP5 pilot was to develop a new vision for Next-Gen aeroengines design, investigating the applicability of breakthrough solutions, merging advanced HPC resources with innovative numerical methods. Two aeronautical use cases, i.e. the combustor and the low-pressure turbine (LPT), have been selected to demonstrate improvement over knowledge and control in critical areas of the engine.

For the low-pressure turbine, the goal was to improve its performance while reducing time-to-market and compliance to current standards and regulations. To achieve this, an innovative AI-based aero methodology has been developed by MORFO Design jointly with Avio Aero, a GE Aerospace company.

In addition, new (high-performance) analytics have been developed being able to analyze terabytes of CFD output data. Thanks to this revised approach, important insights, which are able to discriminate physical phenomena occurring inside turbine passage, can be derived.

Evidence got in first industrial applications have been extremely positive, outlining a great potential for getting outstanding solutions for next-generation aero-engines.

For the combustor use case, a tool named as U-THERM3D has been developed by UNIFI jointly with Avio Aero by means of a multi-physics unsteady approach acting on different layers, ranging from HW/SW integration, combustion modelling and communication streams among separate domains solving (in parallel) fluid, solid and radiative fields. Simulations' accuracy and wall time reduction have been put as main KPIs to be monitored. First evidence gathered from industrial application analysis is very encouraging.

Parallel to the above steps, important synergies and co-design activities have been realized inside the ACROSS team: standard CFD tools have been deeply investigated to reach out more than halved computational time, AI tools as well as HPDA routines have made a leap in quality, passing from old frameworks to newer ones, integrating with the ultimate, even accelerated, hardware solutions.

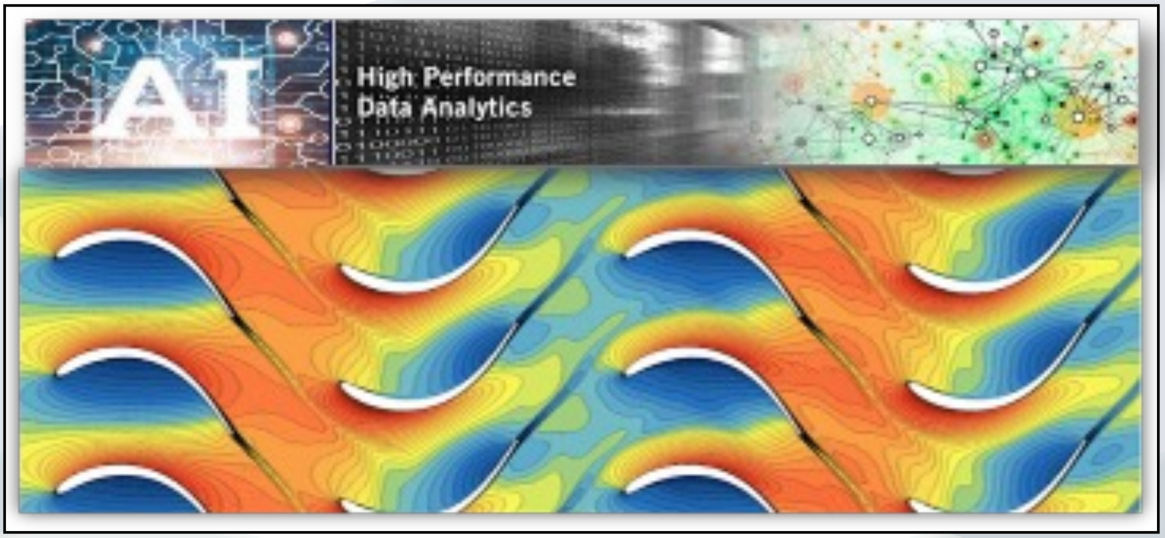

Fig. 5 – A visual representation of the innovations brought by the integration of AI models in the design process of aeronautics components

Weather and Climate Pilot

The WP6 pilot encompassed different workflows aimed at improving numerical processing and data management capabilities, supporting different applications in the Weather and Climate domain. Among these workflows, farming was one demonstrating a big impact, as detailed in the following.

The Farming Use Case

In the ACROSS WP6 farming pilot use case, the improved forecasting skill of a *mesoscale* weather model was examined through the downscaling of new global high resolution (i.e., 4.0km) weather data (HRES) developed by ECMWF. The specific global data have been developed by the ECMWF within the context of ACROSS project. The aim was to provide accurate weather forecasts of 1.0km resolution, covering the Eastern Mediterranean and Greek peninsula.

NEUROPUBLIC used the Weather Research and Forecasting (WRF) system for the above process. Furthermore, the data of NEUROPUBLIC's agrometeorological stations (called *gaiatrons*) were assimilated by WRF's data assimilation component (WRFDA). The abovementioned procedures are computationally ambitious, and they required access to High-Performance Computing systems. ACROSS gave access to both hardware and software infrastructures in order to cover the needs mentioned above. Apart from the access to pre-exascale infrastructures (CINECA Galileo100), installation of WRF and WRFDA, innovations on the Field Data Base (FDB) through exploitation of fast I/O devices for speed up both HRES and WRF, have also been made. The whole procedure is shown in Figure 6. Examples of forecasting products from the abovementioned procedure are shown in Figure 7. Accurate predictions of temperature are very important, especially when they concern cases of heatwaves (Figure 7a) and frosts, which are significantly detrimental to the crops. Furthermore, precipitation patterns become very accurate at a resolution of 1.0km, allowing accurate, reliable and personalized early warnings to the farmers (Figure 7b). The forecasting products also contribute to the estimation of agrometeorological parameters such as evapotranspiration (Figure 7c). The knowledge of a forecasted evapotranspiration is of great importance, as it contributes to a more accurate estimation of irrigation water need and allows a better irrigation planning.

Concerning the performance of WRF on Galileo100, an 84h forecast can be produced in less than 4 hours using 48 nodes (48 CPUs per node). Tests are shown in Figure 8.

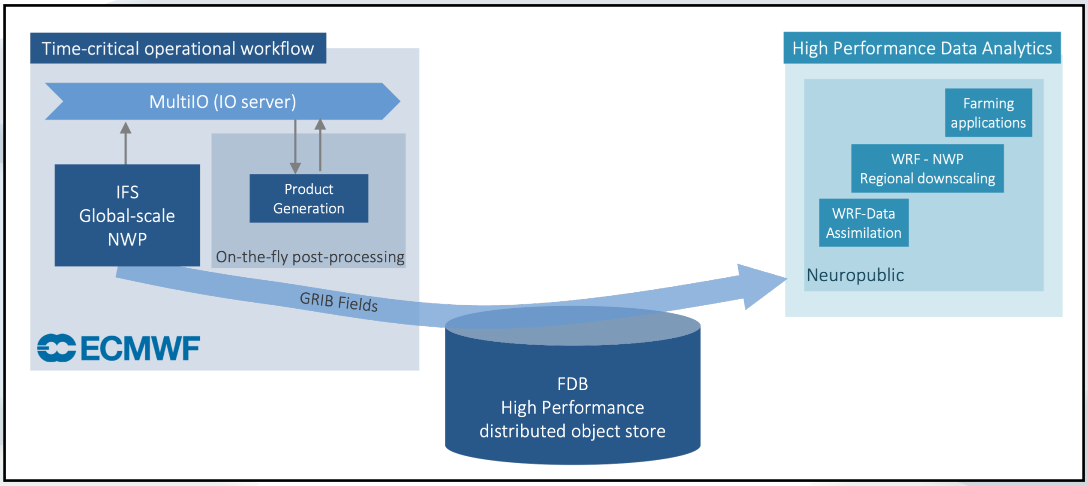

Fig. 6 – Overview of the weather forecasting and farming workflow

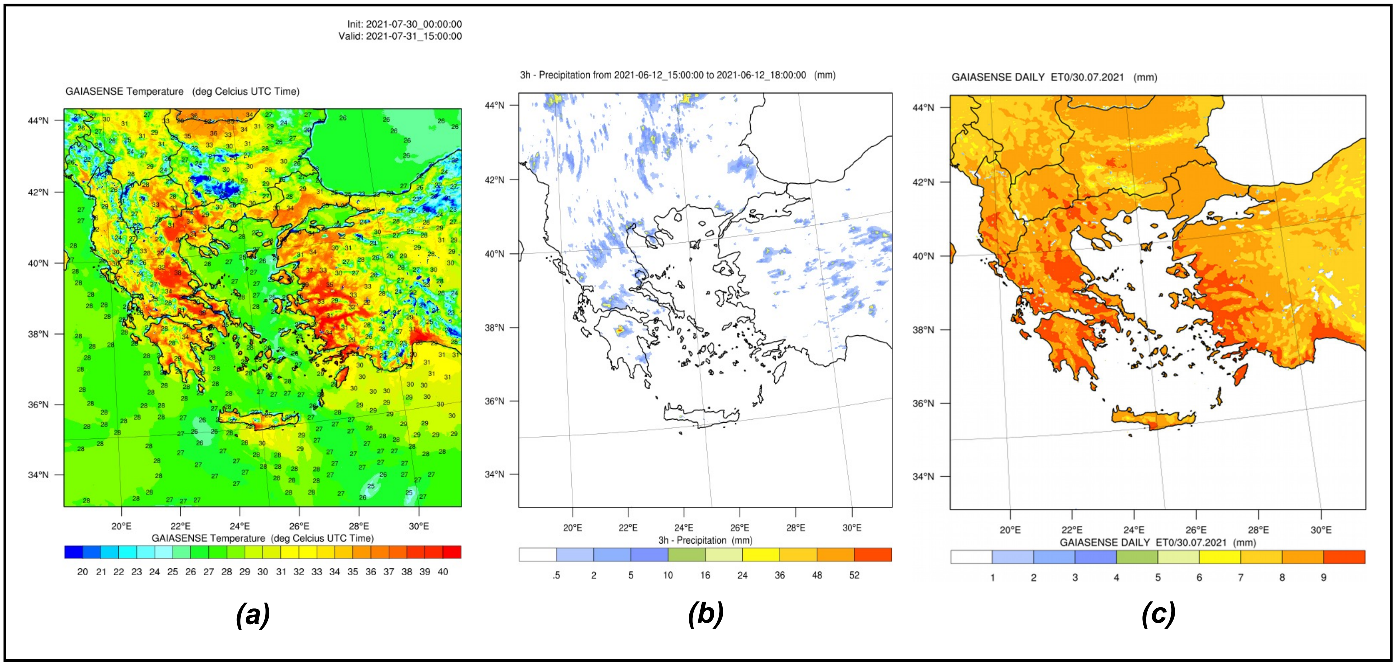

Fig. 7 – Forecast of (a) temperature, (b) precipitation (mm) and evapotranspiration (mm) through downscaling the new 4.0km ECMWF data to a resolution of 1.0km. The dates are not the same for the three pictures

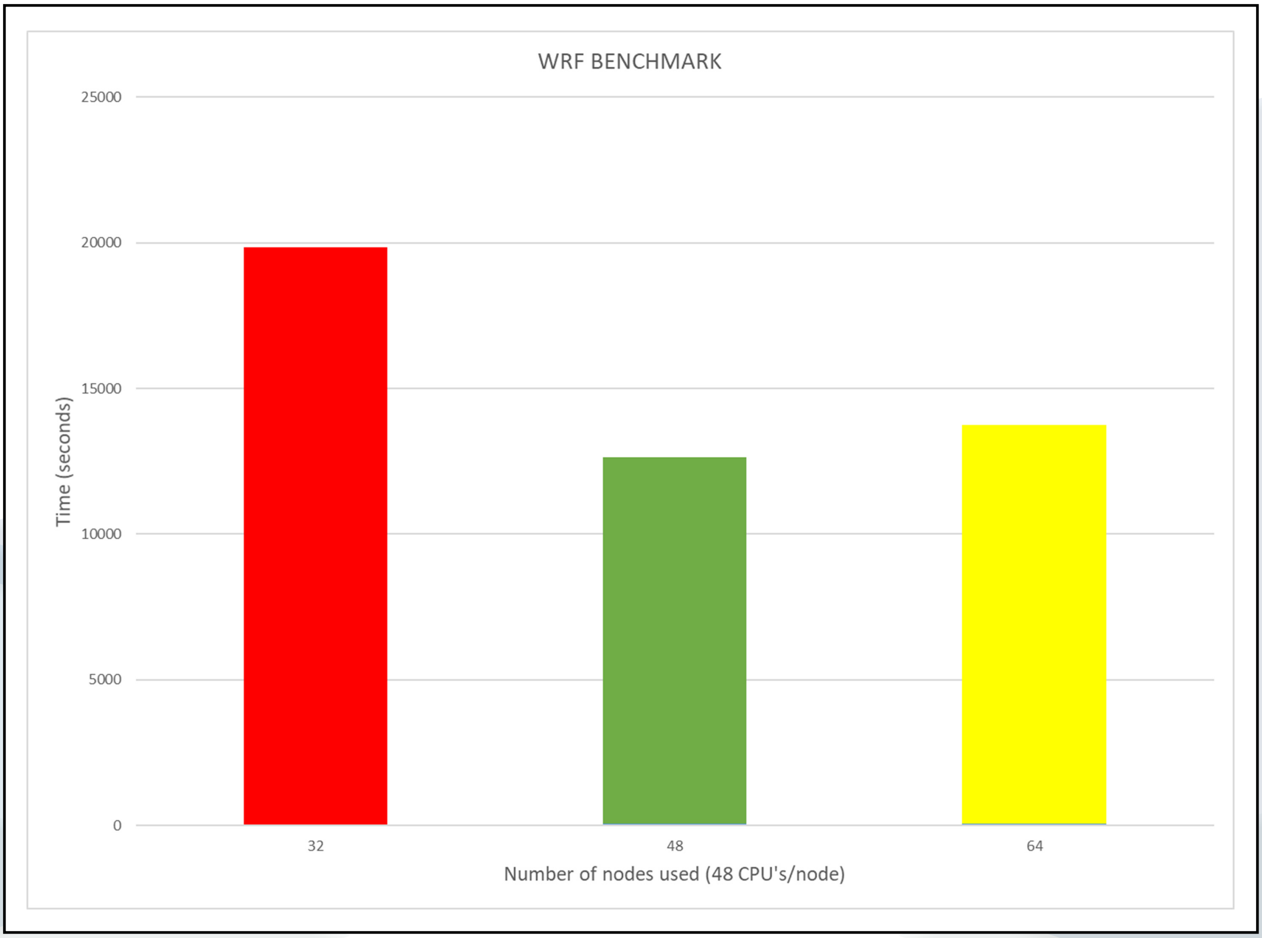

Fig. 8 – Weather Research and Forecasting (WRF) benchmarking

Energy and Carbon Sequestration Pilot

Recent advancements in WP7 have resulted in a GPU-accelerated linear solver beating the CPU-based one, specifically on a test case derived from the SPE11C benchmark. By simulating 2,000 years of CO_2 storage gave a speedup of 1.17 per linear solve on the GPU compared to the CPU on a $3{\cdot}10^6$ cell case. The GPU linear solver uses the newly implemented Diagonal ILU (DILU) preconditioner with graph coloring and row-reordering of the matrix to extract parallelism and access memory efficiently.

The GPU DILU preconditioner code, which was run on an Nvidia RTX 4070TI card, was faster than the more sophisticated constraint pressure residual (CPR) preconditioner, which was run on an AMD Ryzen 5950x using all its 16 cores. The two-stage CPR preconditioner is designed for reservoir simulation and uses algebraic multigrid and ILU0. CPR is a stronger preconditioner than DILU. Despite this, the parallel capability of the GPU surpasses CPU performance when the grids are sufficiently large.

This is an important milestone for OPM Flow, as it demonstrates the usefulness of its GPU solver implementation in the case of large-scale CO_2 storage simulations (Figure 9 shows the performance of the two solvers compared together).

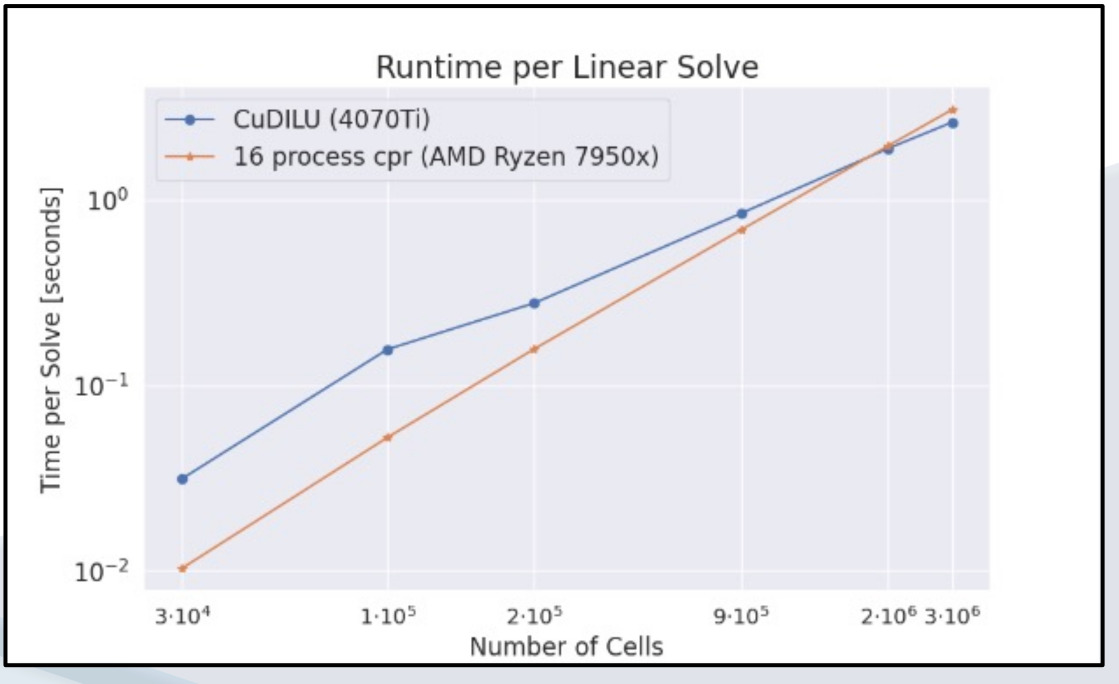

Fig. 9 – Comparison of the solver using a DILU preconditioner running on Nvidia GPU vs. solver using the CPR based preconditioner

Visit our website www.acrossproject.eu to see all our latest developments