ISSN 2599-8293

European
Commission

CORDIS Results Pack on

supercomputing

A thematic collection of innovative EU-funded research results

July 2024

How the EuroHPC JU is accelerating European research and industry

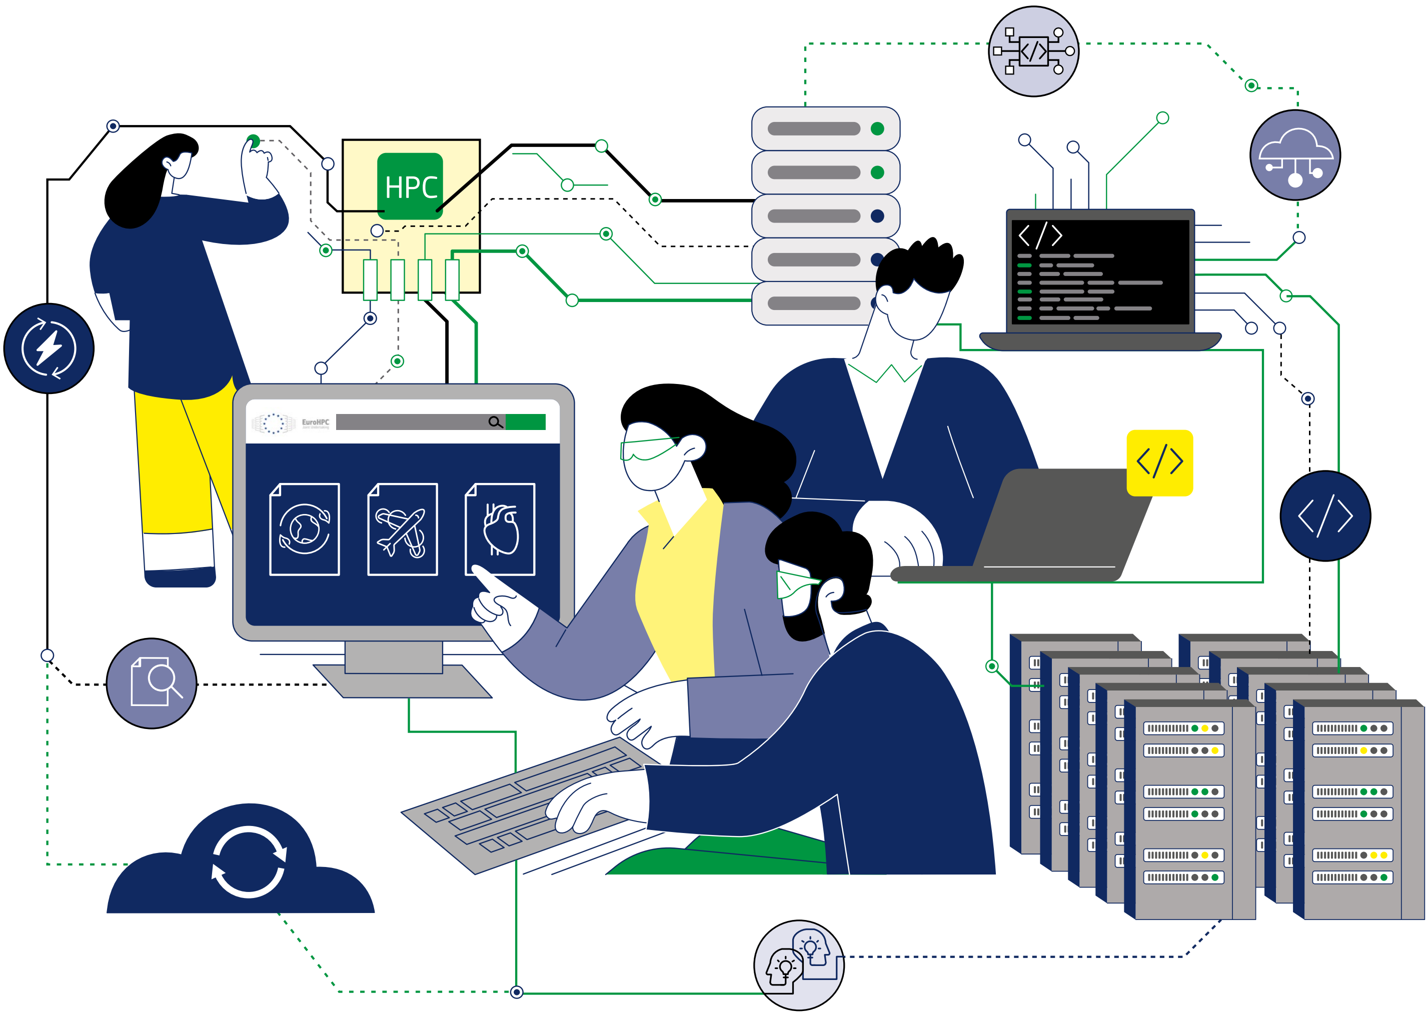

Research and Innovation

Contents

Editorial

How the EuroHPC JU is accelerating European research and industry

Supercomputers can answer some of our most complex scientific questions. The CORDIS Results Pack on supercomputing highlights 15 European projects supported by the unique European High Performance Joint Undertaking (EuroHPC JU), which aims to put Europe at the forefront of the high-performance computing and quantum computing revolution. The investment made by the EuroHPC JU supports the objective of building a world-class supercomputing and quantum computing ecosystem in Europe, one that will accelerate research and industry, boost European competitiveness and innovation and improve the quality of life of European citizens.

Supercomputers are advanced systems with extremely high computational capabilities. They are able to perform calculations which require more power than traditional computers have the capacity for, and can deliver solutions to some of the world's most computationally complex problems.

The high-performance computing (HPC) services offered by supercomputers are essential for discovering new drugs, speeding up the diagnosis and treatment of diseases, modelling the changing climate, anticipating severe weather conditions, increasing cybersecurity, creating more efficient wind turbines and aircraft and even developing more sustainable products.

Yet, developing and building supercomputing systems presents a host of challenges. The EuroHPC JU is investing in the design of computer hardware such as powerful microprocessors and architectures to connect thousands of these individual processing cores as well as the design of software which will allow them to work efficiently in unison. One benefit of this co-design approach is that it will reduce the high power consumption and heat generated by these systems.

The EuroHPC JU is a joint initiative created in 2018. It pools together the resources of the European Union, 35 European countries and three private partners with the ambition of making Europe a world leader in supercomputing and quantum computing.

To this end, the EuroHPC JU is procuring and installing supercomputers across Europe. No matter where in Europe they are located, European scientists and users from the public sector and industry can benefit from these very powerful EuroHPC supercomputers to help them solve scientific problems that involve large data sets. The EuroHPC JU is also in the process of deploying a quantum computing infrastructure offering diverse quantum computing technologies, integrated into different supercomputers across Europe.

In parallel, the EuroHPC JU is funding an ambitious research and innovation programme to develop a full European supercomputing supply chain: from processors and software to applications to be run on these supercomputers and know-how for developing strong European expertise.

Starting in 2024, the EuroHPC JU now has a new artificial intelligence (AI) mandate to develop and operate AI factories. These comprehensive open AI ecosystems centred around EuroHPC supercomputing facilities will support the growth of a highly competitive and innovative AI ecosystem in Europe.

In this Results Pack, you will discover 15 EuroHPC JU projects that have contributed to the technological leadership and digital autonomy of Europe.

How AI can help us tackle future exascale workloads

Supercomputers process supersized amounts of data. The EuroHPC JU-funded ACROSS project combines high-performance computing with artificial intelligence to accelerate research on aeronautics, climate science and carbon capture.

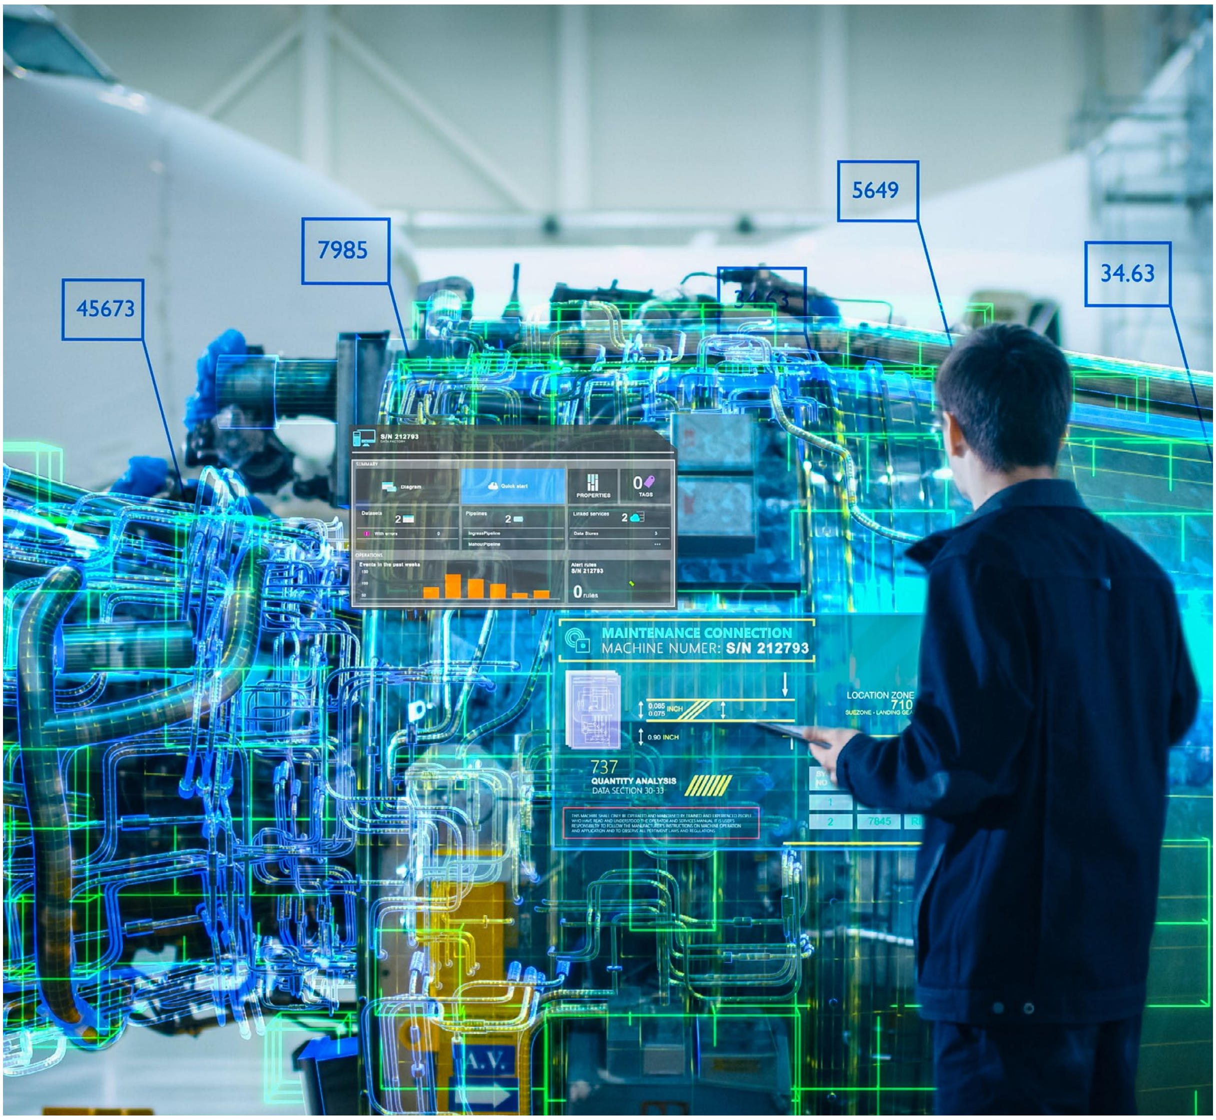

Europe is developing state-of-the-art supercomputers that provide researchers with unparalleled computing power to support scientific discoveries. Yet, complex simulations such as weather forecasts and aeronautical designs also create massive amounts of information that can be hard to analyse using traditional methods.

The ACROSS project uses the artificial intelligence (AI) models to enhance existing high-performance computing (HPC) tools. AI could revolutionise HPC software tools in various technological domains such as speeding up the design process for critical aircraft components, enhancing weather forecasting accuracy, and advancing carbon capture and storage simulations to support a cleaner energy future.

Alberto Scionti, a senior researcher at Fondazione LINKS in Italy and ACROSS dissemination manager, explains: "AI models have progressed so far that they are now able to capture the complexity of our world. This represented an indisputable opportunity to enhance HPC software tools and thus accelerate the ability of researchers to explore and discover more."

AI models have progressed so far they are now able to capture the complexity of our world.

The ACROSS consortium, comprising 17 project partners from Europe, set out to design and develop a convergent platform integrating HPC, big data and AI. The ultimate objective was to create a modular, flexible software stack capable of handling complex workflows on modern supercomputers, combining traditional simulations with AI models and big data processing.

As a result, it will be easier for users to describe their workflows, simplify how they want them to run and make sure that data and tasks flow smoothly even when using cloud resources.

Enhanced scientific progress

ACROSS has enabled researchers to harness supercomputer power more effectively. The project has developed new tools that have allowed for deeper analysis and more complex simulations, accelerating the pace of discovery.

In the field of aeronautics, ACROSS reduced the time needed to run simulations and accelerated the design process for critical aircraft engine components. This will help to deliver more efficient and environmentally friendly aircraft. "The aeronautics pilot advanced the AI model to a new stage, targeting multiple parameters at a time, further increasing the speed of each single workflow run," Scionti says.

ACROSS' achievements also include the optimisation of data processing for weather forecasting. This resulted in more accurate weather predictions, crucial for mitigating the impact of extreme weather events. Finally, the project also improved the simulations of carbon capture and storage, which are essential for reducing greenhouse gas emissions and mitigating climate change.

Advanced computational tech

ACROSS ensured predictable execution times for complex workflows, allowing researchers to experiment with AI models more efficiently. This deterministic workflow execution has enabled scientists to plan and conduct experiments with greater precision and reliability, and has led to improving the overall research process.

Additionally, the exploration of human brain-inspired architectures by this project paved the way for future gains in computing efficiency. These neural networks promise significant improvements in processing speed and energy efficiency, making computers – and research – much faster.

The ACROSS project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

PROJECT
ACROSS – HPC Big Data Artificial Intelligence Cross Stack Platform Towards Exascale

COORDINATED BY
LINKS Foundation in Italy

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955648

PROJECT WEBSITE
acrossproject.eu

Making Europe's supercomputers user-friendly

As high-performance computing systems become increasingly complex, using them efficiently becomes a challenge. This is where the EuroHPC JU-funded DEEP-SEA project stepped in.

Supercomputers are a vital tool for European science and industry. These powerful machines can simulate the Earth's climate, model the human brain, and design innovative products and materials for industry.

A typical high-performance computing (HPC) system combines the central processing units found in all computers with specialised accelerators and a combination of memory and storage options. While such a 'mix and match' approach boosts performance, the inherent complexity makes supercomputers tricky to use. The DEEP-SEA project set out to build and integrate bespoke open-source software to help scientists and businesses get the most out of these powerful machines.

"It is crucial that Europe invests in improving supercomputer software stacks to better support its heterogeneous HPC systems," says Estela Suarez, joint lead of novel system architecture design at the Jülich Research Centre in Germany, and DEEP-SEA project coordinator. "So this was our goal: to deliver an advanced software stack for high-end (pre- and exascale) European HPC systems."

The DEEP-SEA approach

Imagine a computer with a huge number of separate brains and different types of memory. DEEP-SEA wanted to ensure all these parts worked together seamlessly, always solving complex computation tasks in the most resource-efficient way, and always placing data in the best spot to achieve this.

The project consortium included 14 European partners, each bringing their own expertise. By integrating existing tools and combining their functionalities, they were able to conduct a comprehensive development process driven by continuous co-design, with the application developers working on real-world use cases from the European HPC community.

The DEEP-SEA software stack was built around the concept of 'Optimisation Cycles'. These were specific flows in the stack that addressed particular objectives. For instance, the 'modular supercomputing optimisation cycle' assists users in selecting the most appropriate compute elements and mapping their applications onto a modular HPC system accordingly.

By the end of the project, DEEP-SEA had delivered a comprehensive, integrated HPC software stack that included low-level libraries, system software, runtime, programming environments, and tools to improve the performance and usability of supercomputers with mixed components.

On top of that, the project also created documentation and training materials to familiarise end users with the features and benefits of using the DEEP-SEA stack and tools, and to help them to navigate and decide which ones to use depending on their individual applications and objectives.

Surpassing expectations

Collaboration was key, and working with similar projects such as IO-SEA (p. 15) and RED-SEA (p. 27) turned out to be even better than expected, with all three projects working together seamlessly, achieving the integration of their respective results and helping to push each line of development further.

Suarez explains the impact of this collaboration: "The large amount of technical work that came out of this experience was not something we expected. It was an enriching experience for all projects involved."

A software stack is never truly finished.

As a result, DEEP-SEA published over 65 high-quality scientific papers in just three years. This accomplishment is testament to the quality of the research and development delivered by the project.

Software for the future

The project was carried out with support from the EuroHPC JU an initiative set up to develop a world-class supercomputing and quantum computing ecosystem in Europe. Though DEEP-SEA concluded in March 2024, the work is far from over. "A software stack is never truly finished, and hardware advancements will need ongoing modifications and improvements to keep up with the pace of changes generated by HPC research," concludes Suarez.

The real success of the DEEP-SEA project lies in its long-term, participatory approach. In combining existing tools and expertise across the EU, DEEP-SEA unlocked new functionalities for HPC that were previously unavailable, and paved the way for straightforward, efficient use of Europe's supercomputers.

PROJECT
DEEP-SEA – DEEP – Software for Exascale Architectures

COORDINATED BY
Jülich Research Centre in Germany

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955606

PROJECT WEBSITE
deep-projects.eu/project/deep-sea

Improving access to supercomputers to boost meteorology, motor design and disaster reduction

Future supercomputer-enabled innovations could be held back by inadequate workflows. A software solution from the EuroHPC JU-funded eFlows4HPC project is already making waves in the field of disaster risk reduction.

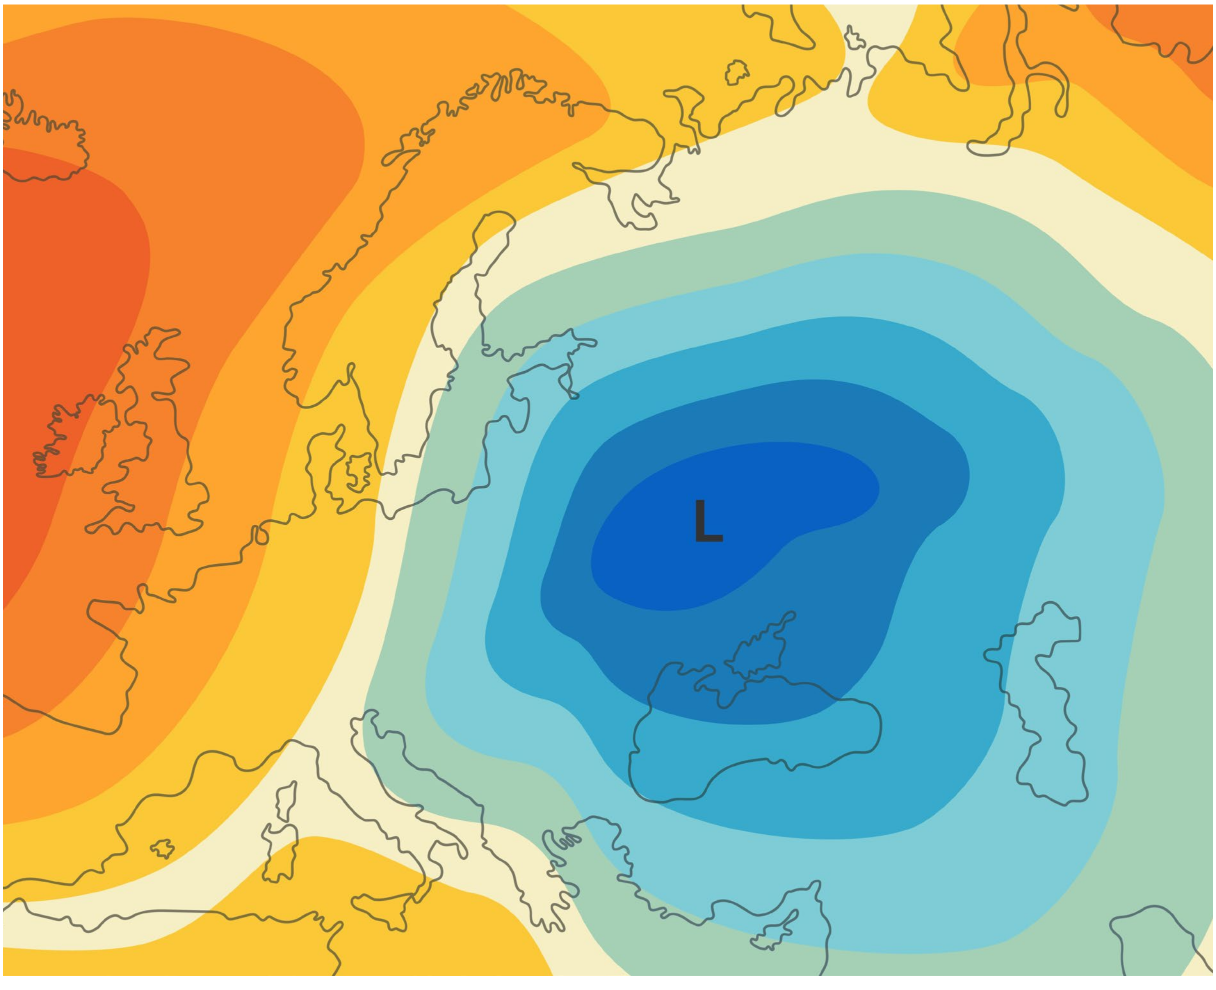

Europe's high-performance computing (HPC) capacity will soon reach exascale speeds, with systems more powerful than 1 million of the fastest laptops. This will allow a wide range of next-generation innovations.

Today's data-intensive tasks, such as weather forecasting, are dependent on computationally heavy workflows combining modelling and simulations and new data analytics along with artificial intelligence (AI) large language models.

Current workflows and associated tools, designed for use in less powerful supercomputers, will probably need to be adapted for these new machines. "The programming models and tools of the applications underlying many current workflows are different, so we need to simplify how they work together," explains Rosa Badia from the Barcelona Supercomputing Center and coordinator of the eFlows4HPC project.

The project has called their approach 'HPC Workflow-as-a-Service' because their software solution, hosted on a user-friendly platform, seamlessly integrates various applications.

"Our approach – which was not possible before – has simplified access to HPC systems, opening them up to non-expert users, spreading the scientific, social and industrial opportunities of supercomputing," says Badia.

So far, over 30 scientific publications have been generated using this approach, alongside a policy white paper on natural hazards.

Under the bonnet

The eFlows4HPC project added new functionalities to existing tools, supporting larger and more complex HPC workflows, and cutting-edge technologies such as AI and big data. At the solution's heart lies a series of software components, collectively known as a software stack, that implements HPC-based workflows, from start to finish.

Firstly, the workflow's various applications – such as simulators, data processing and machine learning predictions – are developed using the PyCOMPSs programming environment.

Data pipelines are also set up to coordinate how workflow data is managed, combining these computational and data management aspects of the workflow.

Finally, everything needed to run the software is encapsulated in a container file and installed on the HPC system. Users can then run their workflows using a dedicated interface.

Testing with diverse user communities

To demonstrate the potential of their approach, the team conducted three use-case trials with very different user communities.

The first use-case trial created is a digital twin of manufacturing processes and assets which simulates experiments with various research and development scenarios cost-effectively. For example, a project partner, CIMNE, collaborated with Siemens to produce a model of an electric motor able to avoid overheating.

As this partnership produced encouraging results, CIMNE now intends to create a spin-off company to exploit these outcomes.

The second focused on Earth system modelling, using two AI-based workflows to increase the efficiency of both Earth system models and prediction models for North Pacific tropical cyclones. These simulations could ultimately help improve global early warning, adaptation and mitigation systems.

With both workflows returning positive results, extensions are now in use within the interTwin project and DestinE initiative.

Lastly, workflows to prioritise access to supercomputing resources known as 'urgent computing' were developed (likely for the first time in Europe) to predict the impact of natural hazards, particularly earthquakes and tsunamis.

Our HPC Workflow-as-a-Service approach has simplified access to HPC systems, spreading the scientific, social and industrial opportunities of supercomputing.

"Thanks to the dynamic flexibility of our solution, extensions of these workflows are now being used in the DT-GEO project, developing digital twins of physical systems to monitor and predict natural hazards," adds Badia.

A springboard for next-generation workflows

The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe. All the open-source software developed by the project, alongside the demonstration workflows, have been made freely available.

"We encourage any user community to access the software and start collaborating to extend its capabilities," notes Jorge Ejarque, eFlows4HPC technical leader.

To ensure the transfer of best practices and results gained from the project, training and workshops were conducted in Barcelona, Munich and Helsinki. The team continues to extend the features and capabilities of all the project's workflows, focused on AI systems and digital twins.

PROJECT
eFlows4HPC – Enabling dynamic and Intelligent workflows in the future EuroHPC ecosystem

COORDINATED BY
Barcelona Supercomputing Center in Spain

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955558

PROJECT WEBSITE
eflows4hpc.eu

Opening the floodgates for supercomputing fluid simulations

By successfully boosting the performance of specialist supercomputing software, the EuroHPC JU-funded exaFOAM project is helping EU industries to maintain a competitive advantage.

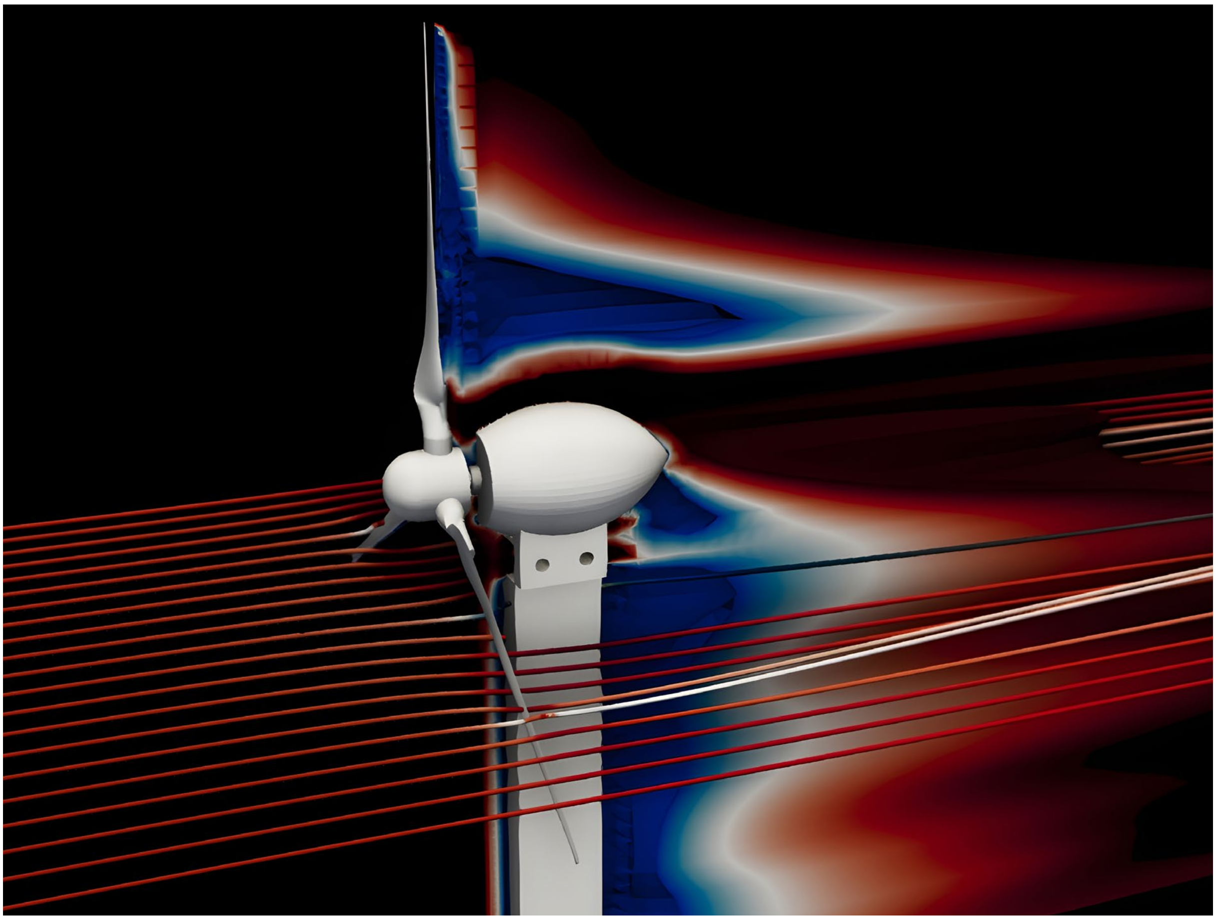

High-performance computing (HPC) enables physical processes to be accurately simulated, avoiding expensive and time-consuming real-world experiments.

A key beneficiary is the field of computational fluid dynamics (CFD), where models predicting the transfer of heat and mass in fluids are used to design applications for sectors, ranging from aerospace to medicine.

"Already a mature technology, future CFD growth depends on more effectively tapping HPC opportunities," explains Fred Mendonça from ESI Group and principal investigator of the exaFOAM project.

The project boosted the performance of OpenFOAM, a highly popular software for CFD workflows, enhancing its ability to leverage the power of HPC throughout the process chain.

By increasing the understanding of areas such as solid mechanics, aerodynamics and heat transfer, OpenFOAM optimises the design of products for ground-vehicle engineering, aerospace, power generation and biomedical applications, amongst others.

"Virtual engineering relies on combining HPC's processing power with the versatility of artificial intelligence techniques. Our work, independent of any particular platform or chip, enables the immense calculations needed to train the algorithms underlying these cutting-edge techniques," adds Mendonça.

Information bottlenecks

During time-consuming computer processing, users rely on receiving intermediate information to assess progress. This data exchange (I/O) can become a bottleneck, especially if the processing core stops until the data exchange is complete.

The project used open-source software, such as ADIOS-2, to maintain efficient I/O file access while running on a large numbers of central processing unit (CPU) cores.

As a result, OpenFOAM successfully ran on HLRS's Hawk supercomputer, where it was deployed to 500 000 cores, for which the project received the HPC Technology Innovation Excellence award.

"To the best of our knowledge, this is the largest number of cores ever deployed for an OpenFOAM simulation," says Mendonça.

Complementing these enhancements, in 2023, OpenFOAM's release software included exaFOAM's profiling tools, which identify processing bottlenecks.

Legacy code

The project also enhanced OpenFOAM's code so it could run on CPU and graphics processing unit (GPU) new chip-architectures.

The main challenge to overcome was that legacy CPU code (such as C++) doesn't necessarily work on more powerful GPUs. The team developed a solution for one part of the CPU's code – known as the linear solver – to be handed over from the CPU to GPU. Work also started on enabling legacy CPU code to be executed directly on GPUs, using industry-standard accelerators (OpenMP).

While the I/O and GPU solutions resulted in increased performance boosts, the team also tweaked some of exaFOAM's algorithms for more incremental gains and to reduce communication bottlenecks across the code. "Combined, these approaches upped OpenFOAM's performance at least tenfold," notes Mendonça.

Industry applications

The project launched several HPC performance benchmarks (Industrial applications and Grand Challenges) and made them publicly available.

"We selected everyday examples, across various sectors, critical to things like: improved equipment design and performance, fuel efficiency for cost-effectiveness, ensuring public safety and comfort and reducing environmental impact," explains Mendonça.

To the best of our knowledge, this is the largest number of cores ever deployed for an OpenFOAM simulation.

Specific industry challenges included: reducing vehicle drag to increase range and optimising turbine placement in windfarm arrays to maximise energy generation. These resulted in 19 benchmark cases.

The project team has now partnered with the EU-funded EXASIM project, which shares the objective of enabling exascale-capable CFD simulation using OpenFOAM.

PROJECT
exaFOAM – Exploitation of Exascale Systems for Open-Source Computational Fluid Dynamics by Mainstream Industry

COORDINATED BY
ESI Group in France

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/956416

PROJECT WEBSITE
exafoam.eu/

Democratising access to high-performance computing

The HEROES software platform matches high-performance computing resources with digital pioneers who could benefit from emerging digital technologies – helping launch a new era of European innovation.

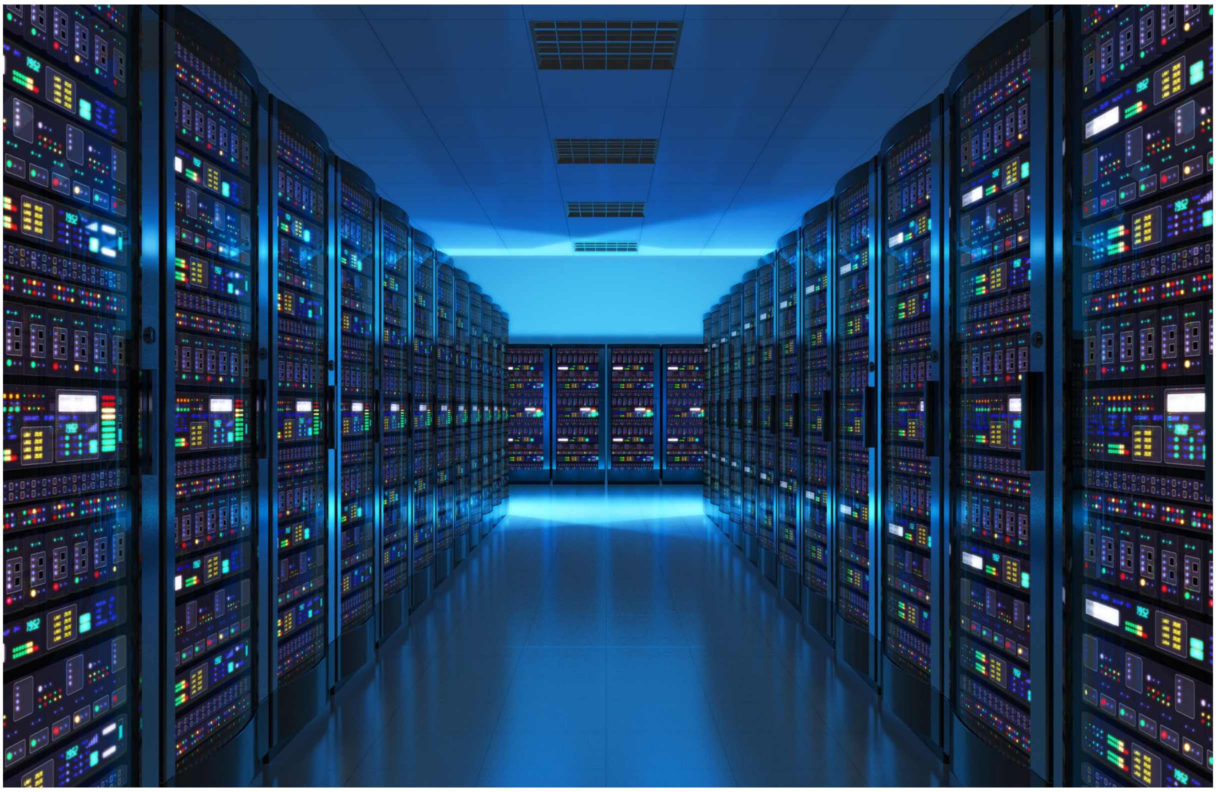

© Scanrail/stock.adobe.com

If industry or scientific communities are to take advantage of fast evolving advances in artificial intelligence (AI) and machine learning (ML), they need access to high-performance computing (HPC) facilities to process large volumes of data.

But these are costly and time-consuming to set up and run, and their energy use raises environmental concerns.

The EU-funded HEROES project has developed software which matches HPC resources to users, to enable them to run simulations and ML workflows based on criteria such as time, budget and energy efficiency.

"We are entering an era of complex digital choices where computational cost per second is vital. Our solution not only helps users navigate the options but democratises HPC, as users don't need specialist IT knowledge to benefit," says Philippe Bricard, HEROES project coordinator from UCIT, the project host.

HEROES developed and tested a demonstrator trialling workflows from the renewable energy and manufacturing sectors.

Simplifying complex infrastructure options

For users of data-heavy technologies such as AI or ML, selecting the best HPC option to run simulations means not only understanding specialist IT jargon, but also weighing up multiple requirements related to budgets, turnaround times, data security and environmental footprint.

The HEROES solution simplifies the process by finding choices for users.

Most projects like this result in proof of concepts. We wanted to go further and actually start implementing.

In the project's prototype, once logged in, users select a workflow or job template, then input their job parameters (including constraints, such as cost). The software then offers a choice between four strategies, linked to HPC platforms. Once selected, the HPC resource runs the job, with the results returned back to the platform.

In some cases the best solution will be hybrid, with some computing resources allocated in-house, others assigned externally, depending on the security restraints and business needs of the user.

"In the future, our platform will be adapted to various contexts, from private organisations with on-premises HPC clusters and some public cloud resources, to more elaborate arrangements with external facilities once authentication issues are resolved," explains Bricard. "Our ambition is to enable clients to use our platform to build and operate their own bespoke HPC marketplace."

From architectural prototype to commercial solutions

The project's prototype comprises a variety of task-specific modules, principally for: workflow and job orchestration; data transfer and storage; pricing and cost management; and energy monitoring and optimisation.

At the solution's core was the decision-making module which communicates with services providers, such as cloud service providers and HPC centres, centralising workflow information and suitable HPC options.

"Most projects like this result in proof of concepts. We wanted to go further and actually start implementing, to start iterating upgrades based on actual use," notes Bricard.

The result was an integration of the decision-making module into the pre-existing data science platform for HPC environments, OKA, improving the latter's effectiveness.

EAR, the system software for data centre energy management, has also now been integrated into the HEROES platform.

"A number of HEROES's components are already available for integration into current computing infrastructures through our commercial venture 'Do it Now'," adds Bricard.

The project's results, along with its partners' expertise, are in prime position to contribute to the EuroHPC Joint Undertaking initiative, set up to develop a world-class supercomputing ecosystem in Europe.

PROJECT
HEROES – Hybrid Eco Responsible Optimized European Solution

COORDINATED BY
UCIT in France

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/956874

PROJECT WEBSITE
heroes-project.eu

Novel data management for exascale supercomputing

Offering a custom-designed storage architecture and on-demand services, the EuroHPC JU-funded IO-SEA project helps facilitate data-intensive applications such as quantum simulations and climate modelling.

© sdecoret

As demand for exascale supercomputers becomes more widespread, operators will have to increase access and workflow capacity, allowing more users to run increasingly diverse and complex applications.

Exascale systems can perform a billion billion calculations per second. Finding a way to manage and store all that data is a significant challenge, as current storage systems reach their limits and operating systems struggle to cope.

"Future applications will not be able to run using current storage paradigms," says Philippe Deniel, head of the Storage Systems Lab at the French Alternative Energies and Atomic Energy Commission.

As coordinator of the IO-SEA project, Deniel has led the development and implementation of a novel software solution that offers long-term storage able to meet increasing data demands.

IO-SEA is one of three SEA projects, along with DEEP-SEA and RED-SEA, set up to develop complementary technologies for a modular European high-performance computing (HPC) architecture.

Future applications will not be able to run using current storage paradigms.

Storage solutions

A key challenge for exascale computing will be the evolution of how calculations are performed. Supercomputers rely on graphics processing units (GPUs), which are designed to break complex problems into thousands of tasks to be performed simultaneously. This means they also require a lot of memory.

Underlying IO-SEA's solution (known as a software stack, as it comprises several components) are innovative uses of hierarchical storage management (HSM), object stores and 'ephemeral' servers.

IO-SEA uses the data storage architecture known as 'object storage', where elements are clustered together, each containing the data, metadata and a unique identifier.

HSM offers a tiered storage approach which automatically identifies the best storage media for the application at hand, whether that be non-volatile memory express (NVMe) such as solid state drives, non-volatile random-access memory (NVRAM), or even tape reels – prized in supercomputing for their low cost and low power requirements. This tiered structure ensures that frequently accessed data is kept on fast media, such as NVMe, with tape acting as a more long-term storage.

"For effective HSM, it's also important to quickly identify files," notes Deniel. "Our advanced monitoring mechanism collects data in a large database, which our artificial intelligence system accesses to make user recommendations, based on their behaviour."

Finally, each storage server is offered on demand, dynamically scheduled to complete a computational job. Operators use a workflow management module to set up simulations, which are then automatically assigned to run on dedicated compute nodes. The results are sent to the storage system, and these servers 'disappear', with the nodes released for the next operation.

Shared resources

Users operate the IO-SEA system using diverse data access middleware such as POSIX, amongst other protocols.

The system was tested in a number of use cases, including electron microscopy, running astrophysics programmes, climatology and Earth system modelling (partnering with DEEP-SEA), quantum physics simulations and large-scale meteorology and weather forecasting.

"Throughout, we demonstrated our solution's capability to offer a paradigm shift from storage as static and unchanging, to being conceived of a process, which is dynamic and shared," adds Deniel.

IO-SEA's solution will be deployed as part of the EUPEX exascale prototype, to be launched within a couple of years. The software has been made freely available on the code-sharing site GitHub.

The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

"Despite being a collection of several products, our solution, co-designed by end users and system developers, introduces an integrated storage stack pointing the way forward for exascale computing," concludes Deniel.

PROJECT
IO-SEA – IO Software for Exascale Architecture

COORDINATED BY
French Alternative Energies and Atomic Energy Commission in France

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955811

PROJECT WEBSITE
iosea-project.eu/

A supercomputing tool to deliver rapid drug discovery

When health crises such as COVID-19 arise, speed is of the essence. The EuroHPC JU-funded LIGATE project accelerates virtual drug design with an artificial intelligence-enhanced supercomputing tool.

The emergence of digital tools involving high-performance computing (HPC) and big data modelling has opened up drug design to non-traditional health industry players such as big tech companies who are challenging the notion that biological systems are too complex to model.

"We are witnessing a shift from models based on data, to data being generated because it is suitable for training models," notes Andrea Beccari, head of discovery platforms at Dompé farmaceutici, the project host.

LIGATE's main goal was to affirm the role supercomputing resources can play in drug discovery, especially crucial in times of acute health crises such as pandemics. The project developed an artificial intelligence (AI) assisted virtual drug screening tool that is adaptable to any hardware architecture, and forms the functional core of the Exscalate drug discovery platform.

Researchers will soon be able to use a network of European computational resources to discover new drugs without having advanced computing expertise.

The Exscalate platform enables researchers to test a library of molecules against a range of targets, to identify possible treatments for threats such as COVID-19. "Initially, trillions of virtual molecules will be screened to find suitable broad-spectrum antiviral drugs. The best candidates will be tested with antiviral assays and structural analysis methods," says Beccari.

The LIGATE tool generates a 3D simulation of molecules of interest, which are used to explore all the molecular conformations (atomic arrangements) possible within the active site of the target protein. The tool selects the most promising based on the molecule's interactions with the protein's amino acids, and an AI tool ranks each based on their predicted efficacy.

Portability and flexibility

The LIGATE solution integrates both open-source and proprietary components. Open-source modules were domain-agnostic to support the entire application's development and large-scale deployment, and domain-specific for the molecular simulation and virtual screening functions.

"LIGATE benefited from the flexibility, efficiency and scalability of community-driven code, complemented by proprietary modules which allow for its commercial exploitation," adds Beccari.

To take advantage of forthcoming exascale computing resources, compatibility with various hardware architectures, especially GPUs, was crucial.

This is particularly true for urgent computing applications – those reliant on the maximum computing power (often using multiple supercomputers simultaneously) to respond to critical threats, such as the COVID-19 pandemic.

But as much of the key hardware is written in proprietary programming languages, the team had to develop a programming model (using royalty-free SYCL) of the main modules, making the software portable across CPUs, GPUs and FPGA cards.

Increasing the accuracy of drug evaluation

A key project aim was to implement the screening tool within the EU-supported LEXIS platform, established to provide greater access to powerful HPC resources. Using a web interface designed to make it even more accessible, the LIGATE tool was demonstrated at the last EuroHPC Summit, when it ran simultaneously on the Karolina, LUMI and LEONARDO supercomputers.

"Researchers will soon be able to use a network of European computational resources to discover new drugs without having advanced computing expertise," explains Beccari, who acted as LIGATE project coordinator.

A predecessor project, EXSCALATE4CoV, used an earlier version of the LIGATE tool to study repurposed drugs for suitability against COVID-19, identifying Raloxifene as a promising treatment. The drug subsequently passed a phase II clinical trial. "Remarkably, during molecular simulations, the platform also found a potentially safer new vaccine, now being preclinically evaluated," adds Beccari.

The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

PROJECT
LIGATE – LIgand Generator and portable drug discovery platform AT Exascale

COORDINATED BY
Dompé farmaceutici in Italy

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/956137

PROJECT WEBSITE
ligateproject.eu

Predicting extreme weather events with artificial intelligence

A groundbreaking initiative by the EuroHPC JU-funded MAELSTROM project created a range of new hardware and software tools to integrate machine learning into climate modelling.

Climate change is expected to reshape weather systems, potentially exacerbating threats such as extreme weather events. Improved weather and climate modelling will help us prepare for these future impacts. Researchers working on the MAELSTROM project have further developed Europe's high-performance computing (HPC) architecture to enhance our ability to predict future climate events.

The MAELSTROM team came together from across Europe to co-design bespoke computer systems to improve energy efficiency, new software able to incorporate machine learning, and new artificial intelligence applications to aid with weather prediction. These technological advances have already strengthened Europe's HPC capabilities and boosted technical capabilities in weather and climate science.

"This project helped the weather and climate community to better understand what machine learning can do to improve weather and climate predictions, how to make the best use of modern supercomputers, and how to enable the HPC industry to provide the best hardware for the application domain," says Peter Dueben, MAELSTROM project coordinator and head of the Earth System Modelling section at the European Centre for Medium-Range Weather Forecasts (ECMWF) in the United Kingdom.

This project helped the community to better understand what machine learning can do to improve weather and climate predictions.

The machine learning applications developed through MAELSTROM are able to transform climate data into more accurate predictions, by improved processing of observations, post-processing of model data and new forecasting abilities.

The project also led to a deeper understanding of which computer system designs are most efficient for machine learning in physical applications. "In particular the application developments have resulted in tools that are already used at the participating weather forecast centres to improve weather predictions on a daily basis," Dueben notes.

Co-designing new hardware and software

The MAELSTROM project is a collaborative effort in which new products and tools were created through a co-design cycle. The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

Weather and climate scientists developed new applications in collaboration with machine learning researchers, while software engineers worked to enhance the software tools available to the scientists developing new applications in training, inference and benchmarking.

Supercomputing scientists then took these applications and software and used them as benchmarks to test new hardware configurations and computer system designs.

"Feedback from hardware developers to the application developers in turn is allowing us to build more efficient applications and to adjust parallelisation, energy consumption and numerical precision during training and inference," explains Dueben.

Setting benchmarks for the machine learning community

The most important results from the project are the products and tools, says Dueben. These range from enhanced machine learning applications for weather and climate science, and meaningful benchmark data sets for the machine learning community which are published online, through to improved software for machine learning.

Revolutionising weather predictions

Machine learning is currently revolutionising how weather predictions are performed, with pure machine learning models outperforming conventional predictions in many aspects.

The project was completed in March 2024, but with major weather centres such as project partner Met Norway developing tools using machine learning, work is ongoing. "I am convinced that all partners will continue their efforts to further improve the applications, software and hardware," concludes Dueben.

PROJECT
MAELSTROM – MAchinE Learning for Scalable meTeORology and cliMate

COORDINATED BY
European Centre for Medium-Range Weather Forecasts in the United Kingdom

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955513

PROJECT WEBSITE
maelstrom-eurohpc.eu

How supercomputers can help defuse ‘cardiac time bombs’

A new exascale simulation platform will allow ultra-fine modelling of heart tissue. Researchers in the EuroHPC JU-funded MICROCARD project hope it can save lives.

© jimmyan8511

Heart disease is the most common cause of mortality in the EU, accounting for a third of all recorded deaths. Around half of those deaths are the result of cardiac arrhythmia, an irregular heartbeat caused by disorders of the heart’s electrical synchronisation system.

There are many sophisticated numerical models of this system, yet in order to model diseased or ageing hearts accurately, they must account for interactions at the cellular level. This takes an enormous amount of computing power, requiring exascale computers (those capable of performing 1 billion billion calculations per second).

In the MICROCARD project, scientists are building the next generation of numerical cardiac electrophysiology models which can represent individual cells and their connections.

"Some phenomena in the heart, in particular the initiation of arrhythmia, depend on events that happen in a single cell, or in the connection between two cells," explains Mark Potse, a researcher affiliated with the Heart Rhythm Disease Institute (LIRYC) in Bordeaux, France, and MICROCARD project coordinator.

"With our simulation platform we will be able to investigate such events and see how they translate into measurable signals, so we can learn how to detect and maybe defuse these 'cardiac time bombs'," he adds.

With our software, it will be possible to turn such reconstructions into models and see how they might behave, and to simulate the signals that would be captured from such tissue.

Building a digital heart

To make the simulator suitable for exascale computers, the MICROCARD team – a collaboration between many experts in mathematics and computer science – implemented the necessary algorithms in the Ginkgo library, which specialises in solving large-scale problems on supercomputers equipped with thousands of graphics processing units.

The researchers updated the simulator code so that computations can take place uninterrupted during data transfer, and to detect hardware and software problems. Finally, the team developed a special-purpose compiler, a piece of software used to translate the equations that represent the dynamics of the cell membrane into code understood by computers.

"These equations are the second largest consumer of energy in our computations, and at exascale the savings this allows are very significant," notes Potse.

The simulation platform itself is still being developed, and advances in individual components are the most tangible results of the project. Many of these are available free and open-source for large user communities.

One important step is improving software that creates 'meshes' – the geometrical descriptions of cardiac tissue. The team is now able to build meshes representing thousands of cardiac cells, and to use them to test components of the software.

"Such software has been developed with the idea that a few million elements is a large mesh," says Potse. "Our ambitions are about a million times larger."

Hearts and minds

The simulation software will be used by project partners and other researchers for cardiology research. Yet the team hopes its reach will extend to similar biological systems such as nerves, muscles, the eye and the brain.

"Recently, neuroscientists have reconstructed a fragment of human brain down to the individual synapses, using serial-section electron microscopy," says Potse. "With our software, it will be possible to turn such reconstructions into models and see how they might behave, and to simulate the signals that would be captured from such tissue."

The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

The team will be able to continue their work in a new project, MICROCARD-2, thanks to the second funding phase provided by the EuroHPC JU. "This will build on the MICROCARD results and bring the simulation platform to actual exascale supercomputers, the first of which is arriving in Europe this year," concludes Potse.

PROJECT
MICROCARD – Numerical modeling of cardiac electrophysiology at the cellular scale

COORDINATED BY
University of Bordeaux in France

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955495

PROJECT WEBSITE
microcard.eu/index-en.html

Supercomputer simulations help sculpt new aircraft designs

By upgrading aerospace modelling tools to harness the power of supercomputers, developers with the EuroHPC JU-funded NextSim project are helping deliver quieter, more efficient aircraft.

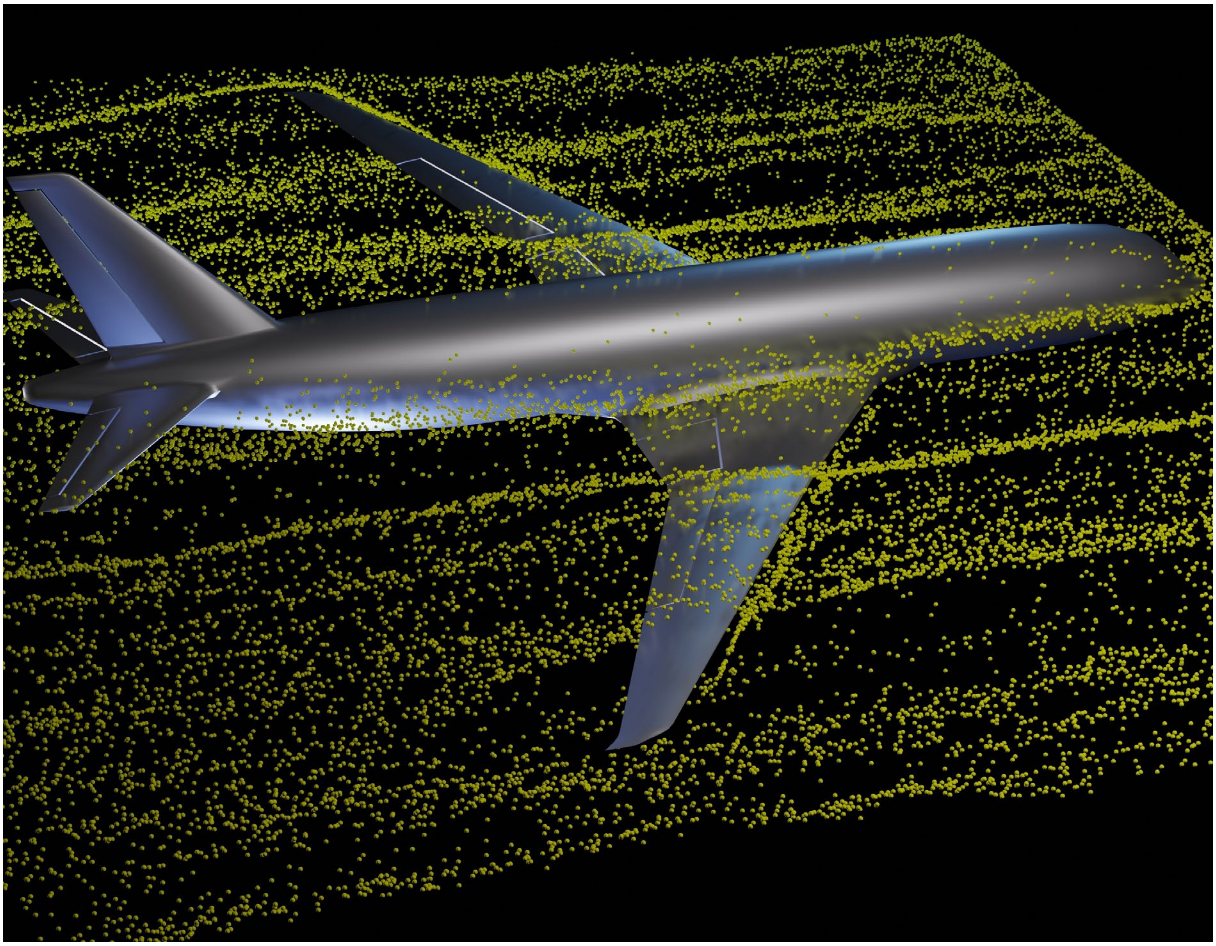

Aviation is a fundamental part of today's society. Growing concerns about the environmental impact of flying means the industry is working to optimise aircraft to make them quieter as well as lighter and more aerodynamic, reducing fuel use and carbon emissions.

Aircraft design is an extremely complex task, requiring the integration of knowledge from multiple disciplines, including aerodynamics, structural design and acoustics. Crucially, investigating new designs often requires the use of expensive wind tunnels and test flights.

There is a clear demand within the European aircraft industry for digital tools that can act as a virtual testing facility. The NextSim project set out to increase current capabilities in computational fluid dynamics (CFD), a critical component of aeronautical design. It achieved this through the development and release of the numerical flow solver CODA (a tool for simulating the behaviour of fluids), which is tailored to run efficiently on Europe's next generation of high-performance computing (HPC) machines.

The outcomes achieved highlight the project's remarkable impact on the aviation industry.

"By re-engineering current CFD tools and leveraging advanced parallel computing platforms, the project aimed to streamline the design process, reduce lead times and enhance the efficiency of virtual mock-ups," explains Oriol Lehmkuhl, large-scale computational fluid dynamics group leader at the Barcelona Supercomputing Center and NextSim project coordinator. "These efforts are aligned with the European Union's goals for aviation sustainability and emissions reduction," he says.

Next-gen simulations

The project was organised around a suite of industrially relevant test cases, covering the aerodynamics of long-range aircraft in cruise flight and in conditions which require high lift (such as take-off and landing), as well as the exhaust flow of jet engines.

These test cases drove the initial planning of the project. All tools and products from the project were also tested in these scenarios to demonstrate their potential.

Simulations developed in the project include analysis of engine exhaust flow for reducing contaminants and noise of aircraft, and the High-Lift Common Research Model, which reveals how air flows over the wings of aircraft during landing and take-off conditions. Thanks to NextSim, these can now be calculated faster and with significantly more accuracy.

Milestone advances

The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

"The outcomes achieved highlight the project's remarkable impact on the aviation industry," adds Lehmkuhl. "CODA now positions itself as a reference-solver for aerodynamic applications within the Airbus group, with substantial implications for the aeronautical market," he notes. "Additionally, the project has addressed market-relevant challenges defined by Airbus, demonstrating the feasibility of solving complex problems with high accuracy and computational efficiency."

The NextSim project also made significant strides in algorithm development, data management and optimisation for advanced HPC platforms. Best practices and guidelines have been shared with industry stakeholders to promote knowledge exchange and adoption of innovative solutions.

Beyond the major advances made in CFD simulations, other achievements of NextSim include the enhancement of current numerical simulation tools for aeronautical design, expansion of HPC usage in the design loop of aeronautical products, improvements in CFD software efficiency, and overcoming existing deficiencies to significantly reduce certification costs by 2050 through virtual design and simulation.

CODA will enter into production through the work of Airbus, DLR and Onera engineers at the end of 2024.

PROJECT
NextSim – CODA: Next generation of industrial aerodynamic simulation code

COORDINATED BY
Barcelona Supercomputing Center in Spain

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/956104

PROJECT WEBSITE
nextsimproject.eu

Programmable accelerators boost supercomputer speeds

By integrating field-programmable gate arrays into supercomputer systems, the EuroHPC JU-funded OPTIMA project developed a range of new applications for industry.

Many state-of-the-art consumer and industrial services, including aeronautics, automobile design, and climate models, depend on emerging high-performance computing (HPC) applications. Meeting the rising demand for such services requires delivering more powerful computers at lower prices.

For some HPC applications, field-programmable gate arrays (FPGAs) are a promising solution. These accelerators can be programmed after manufacturing to suit tailor-made applications. FPGAs can greatly outperform competitive processors such as CPUs or GPUs, granting these significant advantages for HPC applications.

"FPGAs excel at parallel and specialised computations," explains Iakovos Mavroidis, researcher at the Technical University of Crete and OPTIMA project coordinator. "They are particularly effective for parallel, irregular or low-latency tasks in HPC, such as cryptographic algorithms, signal processing and scientific simulations," he says.

OPTIMA is a great example of a small business-driven EuroHPC project.

In the OPTIMA project, Mavroidis and his colleagues used HPC systems integrating FPGAs to develop a series of open-source libraries and industrial applications, such as fluid flow models, with significantly higher energy efficiency.

"Developing use cases on OPTIMA FPGA-based infrastructures was challenging, as we had to balance flexibility and manage resources effectively to ensure good performance without overloading hardware," adds Mavroidis. "However, careful planning, thorough testing, and teamwork among FPGA design, hardware and software teams helped us tackle these challenges."

Integrating FPGA-based accelerators

The OPTIMA team – a consortium of partners from across Europe – started by understanding the requirements of industrial applications and determining the necessary features for the eventual OPTIMA open-source libraries.

Next, the researchers deployed two hardware platforms: one which interconnected four FPGA-based accelerators; the other based on Maxeler's dataflow engines and programming model.

The team then started coding on the FPGA-based platforms, thoroughly testing each aspect, while continuously adding optimisations to the system to make it more efficient. The researchers documented their actions and shared their work with the open-source community.

Demonstrating prototypes

OPTIMA successfully developed four HPC applications on two FPGA-based prototype accelerators. These applications include tools for underground analysis, powerful 'MESHFREE' simulation tools and AI-based robotics simulations – all applications which deal with huge amounts of data and complex tasks.

Additionally, the project created an open-source library on FPGAs. This consists of 31 hardware components which support fundamental linear algebraic operations and computer-aided engineering problem-solving methods, which are crucial for artificial and machine learning applications.

"OPTIMA demonstrated how using FPGA-based technologies can improve HPC systems for industry," says Mavroidis. As the library has been made available to everyone, developers can easily move applications and legacy code to FPGA-supported HPC systems.

The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe. "The OPTIMA project supports EuroHPC JU's goal of making Europe a leader in HPC, and it's a great example of a small business-driven EuroHPC project," notes Mavroidis.

Business acceleration

The team expects the OPTIMA open-source library to continue to grow, to become a valuable resource for software developers aiming to optimise their applications for FPGA-supported HPC platforms.

The expertise gained from adapting applications to the OPTIMA platforms is now being applied by partner SMEs. This includes developing new applications for advanced cloud systems, AI accelerators, GPUs and other chip designs.

PROJECT
OPTIMA – Optimizing Industrial Applications for Heterogeneous HPC systems

COORDINATED BY
Technical University of Crete in Greece

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955739

PROJECT WEBSITE
optima-hpc.eu

Networking tools to boost supercomputer support

New advances in interconnect technologies from the EuroHPC JU-funded RED-SEA project can place Europe at the forefront of exascale supercomputing, bringing benefits across a range of fields, from climate modelling to drug discovery.

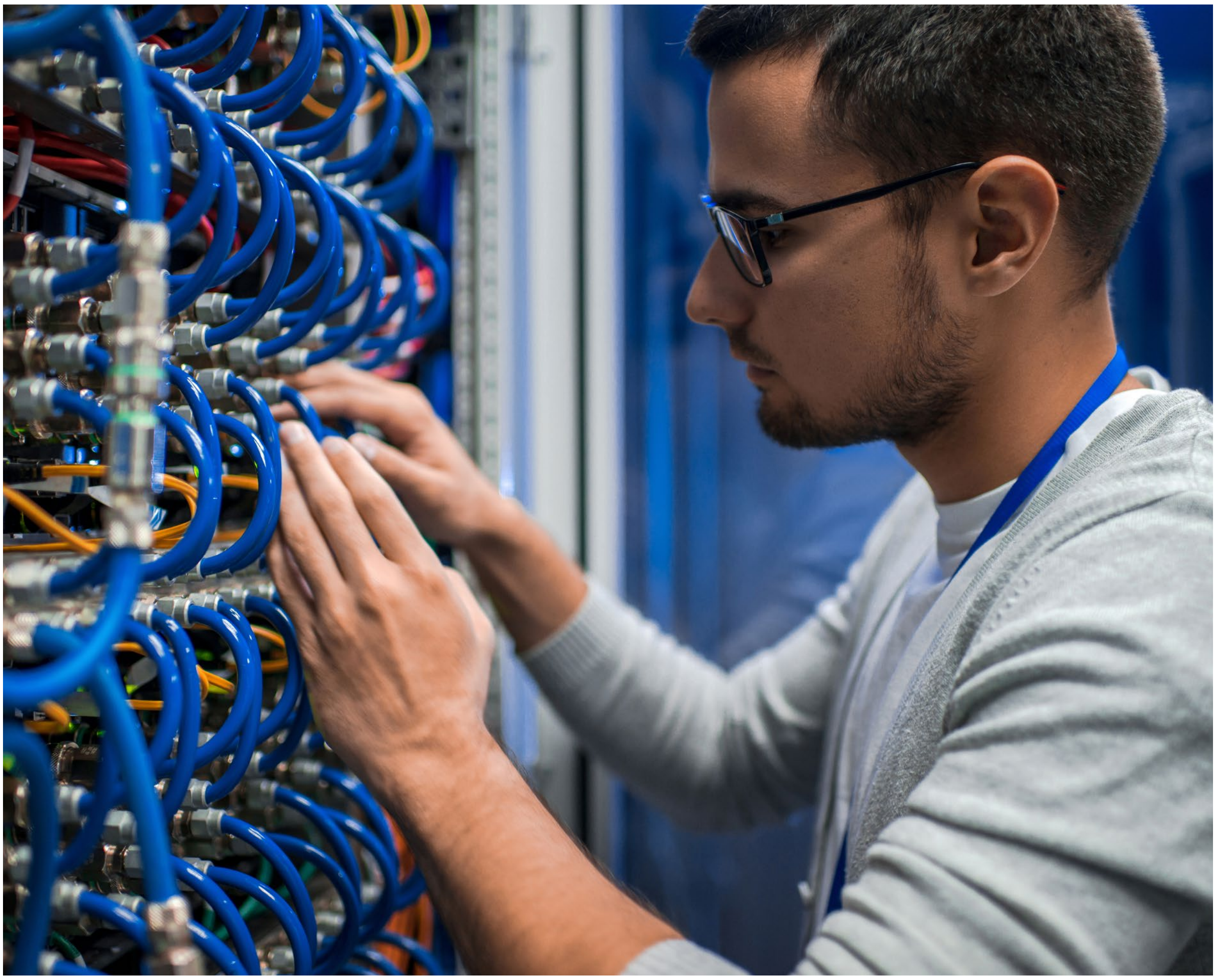

Exascale supercomputers represent a significant leap forward in computing power, capable of performing a quintillion (10^{18}) calculations per second. "The potential benefits are numerous and far-reaching," explains RED-SEA project coordinator Claire Chen, from technology firm Eviden in France.

Applications for supercomputing include climate modelling, astrophysics and genomics. "In essence, exascale supercomputers hold the promise of revolutionising various aspects of science, technology and society by providing unprecedented computational power to tackle some of the most complex and pressing challenges facing humanity."

Connectivity boost

The goal of the RED-SEA project was to lay the groundwork for exascale deployment, by identifying ways of efficiently managing exascale computing systems. Within these systems, interconnection networks serve as the backbone, and play a crucial role in overall performance.

"These networks need to support all the individual connecting nodes, parallel processing systems, efficient connection to the data-centre network, and emerging data-centric and AI-related applications," notes RED-SEA project technical manager Damien Berton.

"Additionally, they need to incorporate features such as efficient network resource management, in-network computing, and power-efficient support for accelerators such as graphic processing units (GPUs)."

The project therefore wanted to develop an innovative, low-latency, scalable and reliable European interconnection network. RED-SEA was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

"The project was built on three pillars," explains Chen. "The first involved leveraging existing technologies, such as BXI, an interconnection network for high-performance computing, as well as the results of previous EU-funded projects, such as ExaNeSt. The second pillar involved exploring innovative new solutions to lay the groundwork for future versions of BXI, while the third pillar sought to develop an ecosystem of users and developers."

Congestion control

Through collaborative partnerships and extensive work, both Chen and Berton believe that significant progress has been made. "We successfully advanced the state of interconnection network technologies," adds Chen.

One significant outcome for example is the advancement of BXI, with a focus on enhancing the current version (BXIv2) and laying the groundwork for its next generation (BXIv3). Another key achievement has been the exploration of new, efficient network resource management schemes.

By harnessing the power of exascale computing, we aspire to create a more prosperous, equitable and sustainable future for all.

"For example, we were able to enhance collective operations, congestion control and adaptive routing," says Berton. "We also extended the European interconnection networks ecosystem through the expansion of the use of BXI networks."

Product range

"We were able to identify and develop 21 RED-SEA exploitable results and achieve two patents," says Berton. Products include an ASIC IP for building an ethernet-integrated circuit, and a software simulator and service for developing highly parallel systems.

Project consortium members will also continue to develop and implement the next generation of BXI. The plan is to integrate BXIv3 within the EUPEX project – a EuroHPC JU-funded exascale pilot initiative.

"Mastering interconnect technologies will provide Europe with an advantage in terms of technological sovereignty and economic return," adds Chen. "We hope that the long-term legacy of this project will include scientific advancement in fields such as healthcare, environmental sustainability and energy efficiency. By harnessing the power of exascale computing, we aspire to create a more prosperous, equitable and sustainable future for all."

PROJECT
RED-SEA – Network Solution for Exascale Architectures

COORDINATED BY
Eviden in France

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/955776

PROJECT WEBSITE
redsea-project.eu

Software solutions for energy-efficient supercomputing

The EuroHPC JU-funded REGALE project developed tools to cut the energy demand of exascale computing. These could benefit European industry and make supercomputing more accessible to researchers.

Powerful exascale supercomputers are expected to prove crucial for next-level scientific research. They are also expected to require a substantial amount of energy.

"Power consumption could be in the order of tens of megawatts," notes REGALE project coordinator Georgios Goumas, from the Institute of Communication and Computer Systems at the National Technical University of Athens in Greece. "Moreover, systems of this scale are quite hard to program, especially when one targets scalability in the order of thousands, even millions of compute nodes." While most supercomputers take power from the central grid, increasing demand could lead to the need for on-site power generation.

Tackling energy efficiency is therefore an essential step to ensuring the feasibility, and widespread adoption, of exascale computing. To address this, the REGALE project brought together supercomputing experts and academics, as well as end users from critical target sectors including renewable energy, enterprise risk assessment and the automotive industry.

Software for supercomputing systems

"We focused on the effective utilisation of computational resources in order to improve application performance, system throughput and energy efficiency," says Goumas. "We also looked at how to ensure the easy and flexible use of supercomputing services by application developers and users."

Several of these tools are now finding their way to industrial or academic supercomputing facilities.

To accomplish this, the project built several prototype supercomputer systems. In particular, the project identified and implemented software, designed to ensure smooth coordination across processes, nodes and systems.

This software included enhancements and coordination of several open-source tools such as OAR, EAR, DCDB, EXAMON, COUNTDOWN, MELISSA, RYAX and others. These are needed for effective resource utilisation and the execution of complex applications.

"These different tools work in concert to support energy-efficient operations at different levels of the architecture," explains Goumas. "They also help to implement a core infrastructure that supports modularity and interoperability, and can integrate any component with minimal modification."

Energy-efficient, sustainable and green solutions

The REGALE solution was successfully pilot-tested with industry partners in order to demonstrate that the tools work smoothly together across a range of different use cases. REGALE supercomputing architecture was used in the design of a car bumper made of carbon-reinforced polymers, as well as the design of hydraulic turbines.

"Several of these tools are now finding their way to industrial or academic supercomputing facilities," adds Goumas. "Project partners have also expressed strong interest in further developing the overall architecture and framework of this solution."

The ultimate aim here is to ensure that the advances made through the REGALE project go on to contribute towards more energy-efficient, sustainable and green supercomputing in Europe. The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

"REGALE started with ambitious goals, a diverse and outstanding consortium of European partners, and a vision to pave the way for greater energy efficiency in the new exascale generation of supercomputers," says Goumas. "Thanks to the commitment and excellent collaboration between all partners, we are now able to present the REGALE toolchain, an open and scalable foundation for sustainable supercomputing."

Energy efficiency at the exascale

This increase in computational capacity has the potential to bring benefits to a range of complex applications including weather prediction, natural disaster prevention and climate modelling. Other possible end uses include personalised medicine and artificial intelligence.

"Beyond these applications though, exascale computing also has the potential to democratise access to supercomputing resources," adds Goumas. "Researchers, academics and SMEs could have easier and cheaper access to advanced computing services."

PROJECT
REGALE – An open architecture to equip next generation HPC applications with exascale capabilities

COORDINATED BY
National Technical University of Athens in Greece

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/956560

PROJECT WEBSITE
regale-project.eu

Applying cutting-edge fluid dynamics at the exascale

By speeding up accurate computational modelling, the EuroHPC JU-funded SCALABLE project could bring huge benefits to industries, including automotive and aerospace.

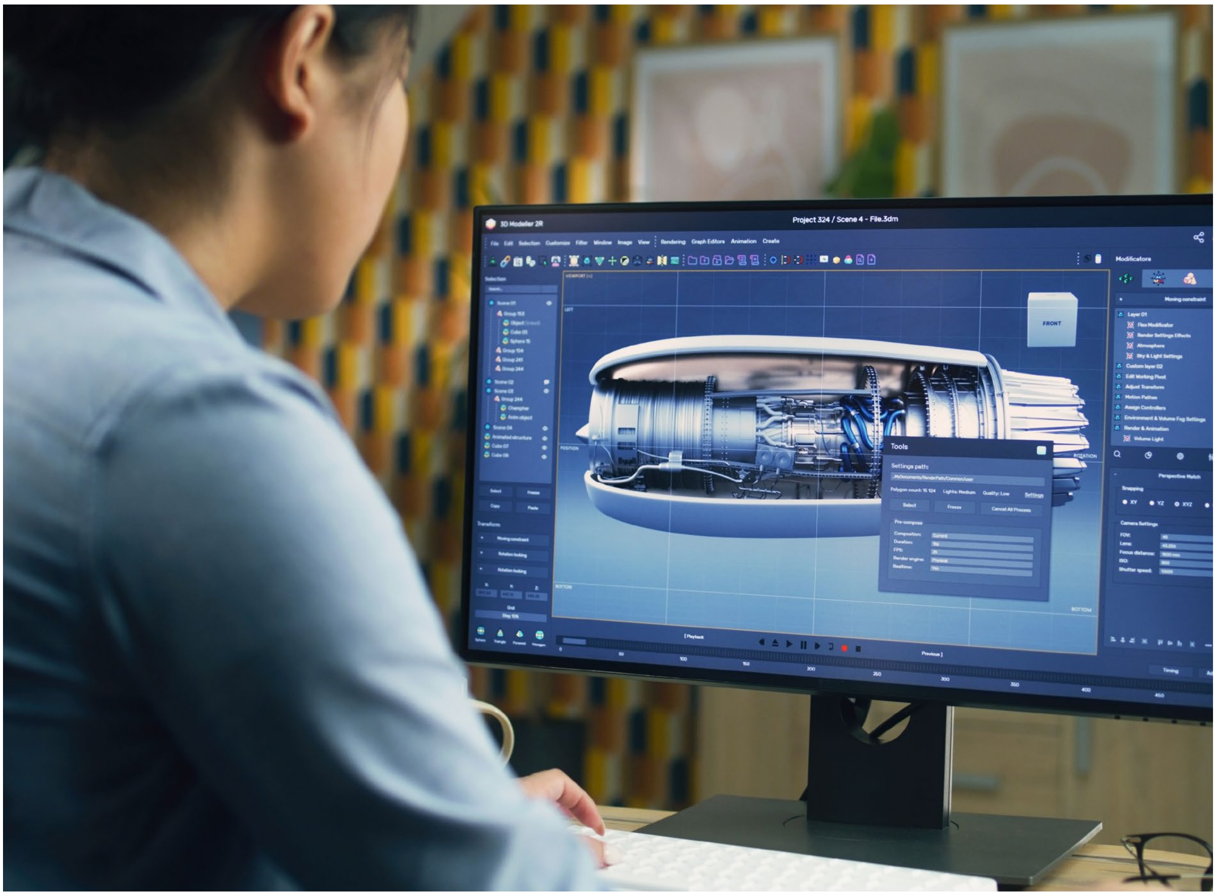

Computational fluid dynamics (CFD) involves the application of numerical analysis and big data to assess and solve problems involving fluid flows. CFD is critical to a wide range of industrial areas, including aerodynamics and aerospace analysis, weather simulation and even visual effects for film and games.

Faster and more powerful CFD calculations would therefore benefit a range of end users, through improving the accuracy and speed of computer models and simulations. To this end, the SCALABLE project, coordinated by CS GROUP in France, sought to develop and test a new class of computational methods, with a view to achieving this goal.

Simulating highly complex fluid flows

To do this, the SCALABLE project team applied what are known as Lattice Boltzmann Methods (LBMs), a CFD approach shown to be significantly faster than conventional CFD approaches. LBMs are numerical techniques that can be used to simulate highly complex fluid flows.

Their strength lies in their ability to easily represent complex physical phenomena, ranging from multiphase flows to chemical interactions. In the context of high-performance computing (HPC), there is huge potential in combining LBMs with advanced supercomputer architectures.

A key drawback however is the fact that LBM techniques consume a lot of energy compared to other numerical methodologies. This has limited their application in the HPC space. The goal of SCALABLE was therefore to find a way of applying LBMs in the context of HPC, to deliver the efficiency and scalability needed for upcoming European exascale computing systems.

Performance, scalability, efficiency

The project brought together a group of eminent industrial and academic partners from three European countries. The objective was to develop LBM-based CFD software, capable of achieving unprecedented performance, scalability and above all energy efficiency.

The project team utilised the public domain research code waLBerla, which offers excellent performance and scalability, and is currently run on some of the largest HPC clusters in the world. The project also made use of industrial CFD software called ProLB (previously known as LaBS). This software solution, based on LBMs, delivers accurate aerodynamic modelling to help engineers make early design decisions. Although at a high level of maturity, there is still room for performance improvements.

Cutting-edge modelling technology

The SCALABLE team successfully combined the benefits of these two existing CFD tools, to deliver cutting-edge modelling technology tailored for exascale computing. In doing so, the project also helped to break down silos between the scientific computing world and that of physical flow modelling.

The prototype software was trialled in a number of test cases involving convection and turbulence modelling. These simulations are critical to the design of things such as airplane engines, wind turbines and cooling systems.

The project also specifically looked at modelling airplane landing gear. This provided a useful test case for assessing the precision and performance of the project's CFD prototype software. The successes and breakthroughs achieved during SCALABLE will now be built upon, by the project's academic and industrial partners.

The SCALABLE project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

PROJECT
SCALABLE – SCAlable LAttice Boltzmann Leaps to Exascale

COORDINATED BY
CS GROUP in France

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/956000

PROJECT WEBSITE
scalable-hpc.eu

Helping next-gen supercomputers keep their cool

New technologies to lower the temperature of computing systems, developed by the EuroHPC JU-funded TEXTAROSSA project, could lead to exascale supercomputers that are faster and more energy-efficient.

While exascale supercomputers promise to accelerate scientific discovery through unprecedented computational power (performing over a billion billion calculations per second), they also promise to consume a colossal amount of energy – and produce a great deal of heat in the process.

“Certain technology gaps still need to be bridged,” explains TEXTAROSSA project coordinator Massimo Celino from the Italian National Agency for New Technologies, Energy and Sustainable Economic Development (ENEA). “These gaps include achieving energy efficiency and thermal control, as well as computational efficiencies. New methods and tools for seamlessly integrating accelerators in high-performance computing multi-node platforms are also needed."

Supercooling supercomputers

Unlike your computer at home, multi-node exascale systems distribute responsibilities across multiple independent nodes, each of which can be a complete computer in itself. This architecture is designed to optimise the performance, scalability and availability of services and applications.

The TEXTAROSSA project has developed a range of tools and technologies that can be integrated into these exascale platforms. The project was carried out with support from the EuroHPC JU, an initiative set up to develop a world-class supercomputing ecosystem in Europe.

"A key technology we developed was an innovative new cooling system," says Celino. Due to the intense concentration of circuity in thousands of central processing units (CPUs), high-performance computing (HPC) systems generate huge amounts of heat. Without adequate measures to remove that heat – such as housing computers in a climate-controlled room – they will become so hot that they eventually fail. These cooling strategies can also be hugely energy-intensive themselves.

While home computers typically use air fans to cool chips, the TEXTAROSSA project went for a more efficient material. "An efficient two-phase cooling technology was developed, where a special fluid is pumped through a closed-loop circuit, which goes around the CPUs," explains Celino.

Working like a refrigerator, the fluid undergoes phase changes, turning from a liquid to a vapour, and then back again, to efficiently absorb the heat from the CPUs and carry it safely away. Since more heat is generated as the computer works harder, the cooling technology was then coupled with an automated thermal control strategy, which controls the intensity of the cooling according to the performance of the CPUs.

Smaller is better

The project team also developed new system software, designed to better manage computational tasks and achieve increased operational efficiencies. "This software offers greater speed in data management and movement, and the seamless management of computational tasks to optimise their distribution over the CPUs," adds Celino. "This enables users to exploit impressive computing power in large parallel runs."

The next step of the project was to develop two different computational node architectures, to be used in future HPC infrastructures. Two prototypes were created using different types of state-of-the-art CPUs, with both exploiting the tools developed through the TEXTAROSSA project.

Several applications, such as exploiting AI and high-performance data analytics, were run on these prototypes. This helped the project team to measure any increased computational and energy efficiencies achieved.

"We wanted to assess whether our techniques could be a useful means of ensuring that supercomputers don't overheat, while still being able to deliver computational capacity for real-life applications," notes Celino.

The challenge now is to further improve these prototypes and work on the miniaturisation and engineering of the technology.

The results confirmed the feasibility of both prototype systems, the efficacy of individual software tools, and the effectiveness of the cooling technology. "The University of Turin asked for a prototype platform to be built, so we actually now have three," says Celino. "The challenge now is to further improve these prototypes and work on the miniaturisation and engineering of the technology."

PROJECT
TEXTAROSSA – Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale

COORDINATED BY
National Agency for New Technologies, Energy and Sustainable Economic Development in Italy

FUNDED UNDER
Horizon 2020-LEIT-ICT

CORDIS FACTSHEET
cordis.europa.eu/project/id/956831

PROJECT WEBSITE
textarossa.eu/

CORDIS Results Pack

Available online in 6 language versions: cordis.europa.eu/article/id/452268

Published
on behalf of the European Commission by CORDIS at the
Publications Office of the European Union
L-2985 Luxembourg
LUXEMBOURG

cordis@publications.europa.eu

Disclaimer
Online project information and links published in the current issue of the CORDIS Results Pack are correct when the publication goes to press. The Publications Office cannot be held responsible for information which is out of date or websites that are no longer live. Neither the Publications Office nor any person acting on its behalf is responsible for the use that may be made of the information contained in this publication or for any errors that may remain in the texts, despite the care taken in preparing them.

The technologies presented in this publication may be covered by intellectual property rights.

This Results Pack is a collaboration between CORDIS and the European High Performance Computing Joint Undertaking.

@HorizonEU
@EuroHPC_JU

@EUScienceInnov

@EuroHPC JU

Print	ISBN 978-92-78-44205-7	ISSN 2599-8285	doi:10.2830/143536	ZZ-AK-24-016-EN-C
HTML	ISBN 978-92-78-44206-4	ISSN 2599-8293	doi:10.2830/895929	ZZ-AK-24-016-EN-Q
PDF	ISBN 978-92-78-44207-1	ISSN 2599-8293	doi:10.2830/98733	ZZ-AK-24-016-EN-N

Luxembourg: Publications Office of the European Union, 2024

RESULTS PACK ON IN-SPACE ELECTRIC PROPULSION

Europe is well on the way to providing its booming satellite market with a cost-effective, highly efficient low-energy electric propulsion system. The CORDIS Results Pack on In-Space Electric Propulsion highlights the main achievements of a cluster of EU-funded research projects working to achieve this breakthrough.

Check out the Pack here:
cordis.europa.eu/article/id/451366

Follow us on social media too!
facebook.com/EUresearchResults
x.com/CORDIS_EU
youtube.com/CORDISdotEU
instagram.com/eu_science

Publications Office of the European Union

EN