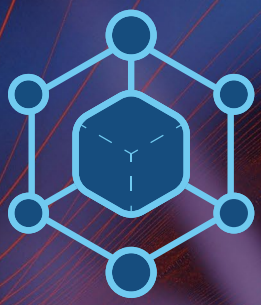

INNO4SCALE

Innovative Algorithms for Applications on European Exascale Supercomputers

Success Story Highlights

Innovative Algorithms for Applications on
European Exascale Supercomputers

Introduction

The EuroHPC Joint Undertaking (JU) is jointly funded by its members with a budget of around €7 billion for the period 2021-2027. A large part of that funding is dedicated to the procurement of a European network of large- to extreme-scale supercomputers, including two exascale systems scheduled to become operational in the near future; at the time of writing, the JUPITER system hosted at the Jülich Supercomputing Centre is ranked 4th in the Top500 List of the world's fastest supercomputers. In addition, the JU's research and innovation programme complements investments in computing infrastructure by supporting the development of high-performance computing (HPC) technologies for future extreme-scale computing systems and for the operation / exploitation of current systems.

In parallel, the JU has funded research in application software used to tackle societal, scientific, economic or environmental challenges. However, exploiting ever faster computer systems poses critical challenges for many of the existing application software suites, with a number of issues that include: the widening gap between sheer compute power and the lagging ability to move data sufficiently fast; finding sufficient parallelism in existing algorithms (which were often initially developed in an earlier era of HPC technologies); and mapping the computations onto heterogeneous hardware, such as multi-core CPUs combined with graphics processing units (GPUs).

Thus, ideas for new algorithms and computational methodologies appropriate for the evolving extreme-scale computing systems are urgently needed – and they are there, in the minds of our researchers. Yet developing and trying out new approaches comes with risks that are often too high to be addressed in the framework of a typical R&D project.

This is where Inno4scale comes in: In an open call for "innovation studies", the project invited proposals for potentially disruptive ideas to reach new algorithmic heights in HPC. The 22 selected studies ran from February 2024 until the end of March 2025, working on bold ideas towards proof-of-concept implementations that reveal the merits of new approaches. The topics include using machine learning in novel ways to speed up repetitive computations, exploiting special structures in linear systems, using time as an additional dimension for improving parallelism, reducing accuracy where it is not needed, and investigating novel ways to organise and distribute computational tasks to fully utilise all available hardware.

Because of the high level of innovation, it goes without saying that not all of the initial plans worked out in full, yet there are remarkable results, many of which will have impact beyond the initial domains investigated in the studies. In this brochure, we do not describe every study in detail, but pick examples illustrating highlights of what was achieved and where the journey could continue regarding future improvements in HPC algorithms.

Several Inno4scale innovation studies tackled numerical linear algebra – both sparse and dense linear systems – which is a central component of an enormous range of HPC applications. The innovation study AceAMG showed how to reduce computational expense and communication demands, in particular synchronisation, speeding up an algebraic multigrid by up to 8×. CBM4Scale focused on a memory-efficient data structure for binary sparse matrices (e.g. for representing large graphs), enabling an over 11× reduction in memory size and up to 5× reduction in time required for multiplication with the associated new matrix multiplication kernel.

Computational mechanics, and computational fluid dynamics in particular, could be said to be a "classic" application field for HPC. The studies we have highlighted here demonstrate the range of very different applications and high-impact innovations: Ex3S employs mixed-precision approaches on GPU systems to enable new possibilities for high-fidelity simulation of hydrogen combustion, targeting clean energy production. Asynchronous multi-GPU algorithms developed in SCALE-TRACK will allow scientists to gain new insights for understanding our environment (specifically cloud formation and development) with unprecedented speed and energy efficiency. ExaSIMPLE creates new performance potential and enhanced robustness for general purpose CFD by including machine-learning methods into the core simulation methodology.

A technique of growing interest is uncertainty quantification (UQ). In many real-world situations, crucial data is noisy, incomplete and / or of limited accuracy. In other cases, essential data is simply unknown and needs to be estimated using simulation. ScalaMIDA addressed such a situation by making algorithmic improvements in software for Bayesian inference, providing a scalable approach for reducing computational costs in earthquake simulation. In MLMC-PinT4Data, a combination of a multilevel Monte-Carlo algorithm with a parallel-in-time approach achieved over 90% of cost savings for UQ in an industrial use case. Both studies developed approaches that could also be transferred to simulations in other domains.

Particle physics is best known for its large-scale and very expensive experimental setups (think large colliders). Can one replace some of this by computing? Yes – but it is a very demanding application. MG4ML showed how to overcome some blockers in this area, using optimised domain decomposition and multigrid techniques. Among several studies that addressed computationally demanding atom-level simulations for materials science, XSCALE was chosen for its up to 100× speedup of certain types of simulation for Generalised Hamiltonian Dynamics. By expanding the size of molecular systems that can be simulated feasibly, the approach will help bridge the gap between computational and experimental data, with applications relevant for studying metal-ion battery chemistry and catalysts for synthetic fuel production.

Finally, the typical sequential workflow of preprocessing-simulation-postprocessing hits feasibility limits at exascale. NEOSC shows how to extract data co-processing tasks from a running simulation on a GPU and map them onto an otherwise idling CPU, thus optimising overall system usage while at the same time cutting time-to-solution in half.

There is much more to discover about the Inno4scale innovation studies and their results than what we can show here. We encourage you to visit Inno4scale's web site (www.inno4scale.eu) to learn more.

Because Inno4scale was conceived to support research at the very frontiers of computational research, most of the work presented here is still "in progress." In the future, it will be further refined, applied to new and wider domains, and combined with other methods. Considering that many of the approaches we present here complement one another, one can expect more gains by clever combinations of the ideas. Furthermore, as with all good research, at the same time that work in the 22 studies has answered some questions, it also raises new ones waiting to be tackled.

AceAMG

Modernising CFD solvers using mixed precision, low communication strategies

Highlights

- New mixed precision, low communication distributed algebraic multigrid (AMG) solvers were created for the Ginkgo software library and integrated into the nekRS software stack. This improved the performance of nekRS in a computationally challenging simulation by up to 8×.
- nekRS with AMG integration is exascale-ready in terms of weak and strong scalability on the Grace-Hopper hardware used in Europe's first exascale system.

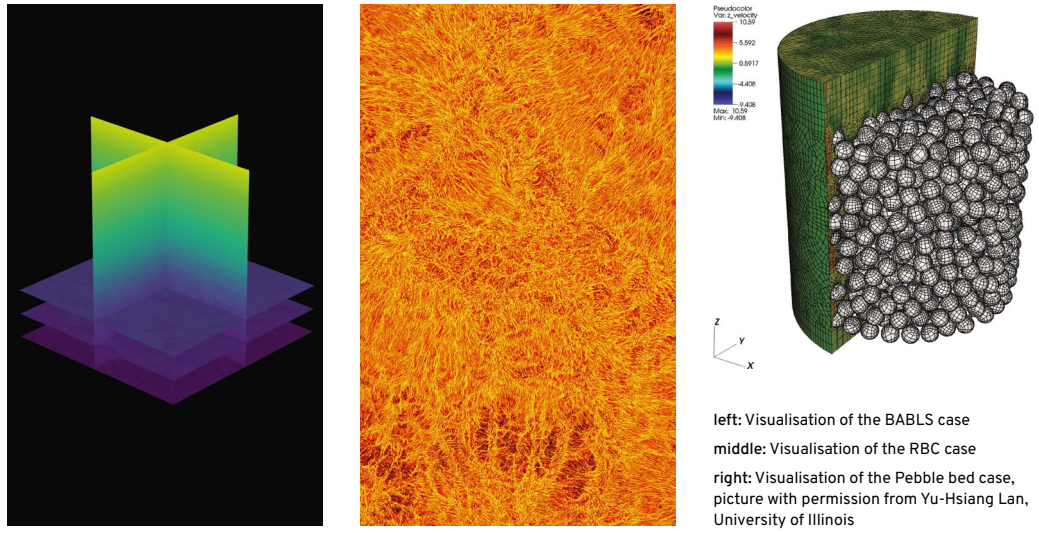

left: Visualisation of the BABLS case

middle: Visualisation of the RBC case

right: Visualisation of the Pebble bed case, picture with permission from Yu-Hsiang Lan, University of Illinois

Challenge

The solution of large systems of linear equations is a key, computationally demanding component of many CFD (computational fluid dynamics) codes. This is in part because the iterative linear solvers commonly used for these calculations are not well-suited to modern hardware architectures. State-of-the-art iterative solvers typically consist of memory-bound, sparse, high-precision operations. Today's most powerful HPC systems, however, increasingly use graphic processing units (GPUs), which are good at low-precision computation but perform poorly when running memory-intensive applications. Because upcoming exascale supercomputers in Europe will contain thousands of GPUs, leveraging these systems efficiently in production runs of CFD simulations requires that simulation software stacks map well to these new accelerated architectures, with CFD solvers needing to become more lightweight and less memory-bound. Strategies for doing so include using mixed precision numerical approaches and parallel programming strategies that make more efficient use of available computing resources.

Solution

The AceAMG innovation study addressed this challenge in two ways. Firstly, it implemented mixed precision functionality into Ginkgo, a software library for linear algebra on GPU-accelerated high-performance computing systems. In addition, the team used compression of message sizes and utilisation of asynchronous MPI communication. These two approaches were combined to create a new mixed precision, low communication distributed algebraic multigrid (AMG) in the Ginkgo library. The team integrated this new solver as an optional backend in nekRS, a software stack for CFD applications. Doing so now gives nekRS users access to solvers and preconditioners for low-synchronisation, low-communication, mixed precision numerical methods.

Benefits & Impact

One key practical goal of AceAMG was to to increase the performance of nekRS on the upcoming exascale supercomputer JUPITER. Using JURECA and JEDI (prototypes of JUPITER, which include NVIDIA GPUs), the team tested the integration of the Ginkgo-based AMG solver within nekRS in several simulation cases: a GEWEX Atmospheric Boundary Layer Study (GABLS), Rayleigh-Bénard convection, and pebble bed cases. Mixed-precision numerical approaches were shown capable of improving performance without decreasing overall simulation precision. For challenging problems like pebble bed cases, the AMG method shows promising convergence improvement — nekRS with Ginkgo's new AMG method outperformed the native nekRS solver by up to 8×. The resulting scientific paper demonstrating the performance benefits is included in the proceedings of the Parallel CFD International Conference (ParCFD) 2024. With implications for simulation in fields such as renewable energy, climate science, and aeronautics, the results of AceAMG will reduce runtime and resource consumption of nekRS production runs. Furthermore, these new methods can be applied to an even broader array of domains in which codes will benefit from Ginkgo's performance portability.

Next Steps

Because many software packages, such as Neko and nekCRF, use or follow similar designs to that of nekRS, the technology developed in AceAMG will be used to improve their performance. The results of this innovation study not only advance Ginkgo's performance in a wider range of real-world applications but also show that this integrative approach can be scaled up for larger simulations.

Organisations Involved
Computational Mathematics (CoMa) Research Group, Technical University Munich; Jülich Supercomputing Centre (JSC).

Contact: Hartwig Anzt, Technical University of Munich, hartwig.anzt@tum.de

CBM4scale

Improved Data Compression for Graph Neural Networks

Highlights

- A new data format called compressed binary matrix (CBM) reduces the number of elements within a matrix while retaining its overall structure, reducing the memory footprint of real-world graphs by up to 11×.
- A new matrix multiplication kernel based on the CBM format enables a 5× speedup over other state-of-the-art sparse-dense multiplication kernels, and performance gains of up to 3× for training and inference phases in graph neural network applications.
- The approach is usable on CPU and GPU architectures, is supported for PyTorch, and is scalable to support large-scale network analysis and artificial intelligence applications.

Challenge

Graph neural networks (GNNs) are a key technology for deep learning based on graph-structured datasets, and are expected to be used widely in future artificial intelligence applications on large-scale HPC systems. To generate GNNs, algorithms must perform long sequences of sparse-dense matrix multiplications, which identify and characterise the significance of relationships between nodes in graphs. Existing sparse matrix formats like coordinate list (COO) and compressed sparse row (CSR) improve the speed of these applications using a Boolean approach in which only non-zero entries are stored; i.e., data points that identify and characterise a relationship between nodes are kept, while data points showing no relationship are deleted. As datasets grow ever larger, though, more powerful data compression approaches are needed. To be compatible with future exascale systems, the solution must run efficiently on CPU and GPU platforms as well as shared-memory architectures, and support parallel computing techniques including multithreading and SIMD (single instruction, multiple data). Moreover, it should be fully compatible with the widely used machine learning library PyTorch and its Autograd engine. It should also be transparently usable with models based on graph convolutional networks (GCN), GraphSAGE, and graph isomorphism networks (GIN) during both training and inference.

Solution

The innovation study CBM4scale set out to improve on existing compression methods by developing a new data format called compressed binary matrix (CBM), which can identify similarities between rows or columns in a matrix. This approach is particularly well suited, for example, for accelerating calculations involving the adjacency matrix of social and collaboration networks, where neighbouring nodes often share very similar neighbourhoods. CBM encodes each row in terms of its differences from another similar row within the same matrix. An algorithm stores only elements that differ between one row and another similar row, resulting in a far smaller number of elements than in other data formats. This approach is based on a minimal row-dependency tree, also known as an arborescence, in which the number of dependencies within a system is minimised while relationships are retained, with dependencies being represented in a branching structure. A matrix multiplication kernel (CBMM) associated with the CBM format performs sparse-dense matrix multiplication by traversing this tree structure, reusing results from previous computations for related rows wherever possible. This allows the product to be computed using fewer scalar operations than in the original matrix. The structure of the dependency tree also enables parallelism, which is exploited using OpenMP and CUDA streams for efficient execution on CPUs and GPUs.

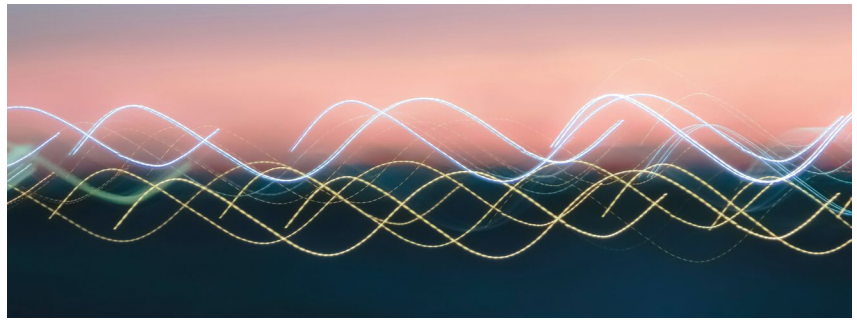

Benefits & Impact

The CBM format and CBMM kernel enable significant improvements in the efficiency of sparse-dense matrix operations without sacrificing accuracy. Experimental results showed that the compressed binary matrix format can reduce the memory footprint of real-world graphs by up to 11×. In addition, parallel matrix multiplication using CBM is more than 5× faster than state-of-the-art sparse-dense multiplication kernels. In the context of GNN applications, the use of CBMM led to performance gains of up to 3× for both training and inference phases, showing strong potential for improving time- and energy-to-solution in real-world AI workloads. The format is inherently parallelisable and future-proof, as further optimisations in underlying libraries will also benefit CBMM. These accomplishments hold the potential to accelerate a variety of scientific and industrial applications. GNNs are essential, for example, for social network analysis, power grid analysis, and prediction of drug-target interactions in computational biology, as well as in AI applications such as recommender systems and natural language processing. The efficiency gains enabled by CBM / CBMM make these workloads more scalable and energy-efficient, and will facilitate their deployment on extreme-scale HPC infrastructures.

Next Steps

Future work will focus on integrating and evaluating the compressed binary matrix format in the context of different graph neural network architectures, including the training stage of these networks. Additionally, the investigators aim to implement and evaluate the CBM format and corresponding multiplication kernels on GPU architectures.

Organisations Involved
Laboratory of Information and Decision Support Systems, Institute of Systems and Computer Engineering: Research and Development in Lisbon (INESC ID); Scientific Computing Group, University of Vienna.

Contact: Alexandre Francisco, INESC-ID Lisboa, aplf@tecnico.ulisboa.pt

Ex3S

Extreme-scale combustion simulation for the clean energy transition

Highlights

- GPU-optimised numerical components enable a spectral element solver for high-fidelity simulation of hydrogen combustion with detailed chemistry.
- The new solver shows high, efficient scalability, addressing the computational needs of direct numerical simulation when deployed on extreme-scale computing systems.

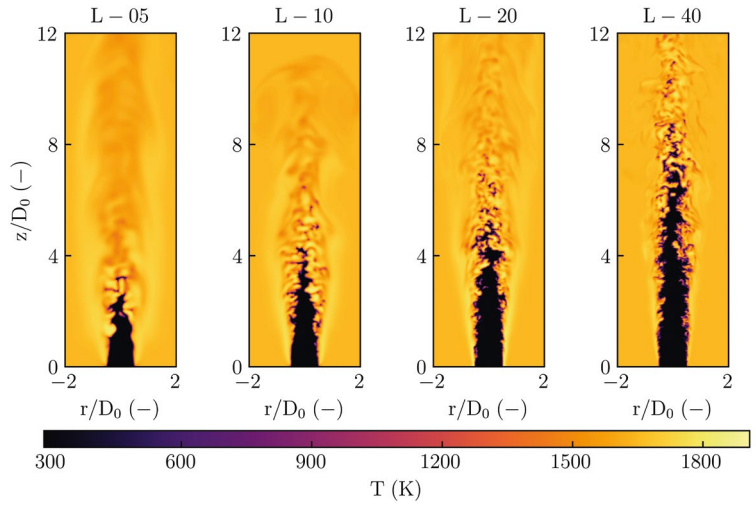

Application of NekCRF to a series of turbulent premixed H_2 / air jet flames. Shown are the temperature contours at various Reynolds numbers (Re = 5000, 10000, 20000, 40000 from left to right).

Challenge

The transition to carbon-free energy systems presents a major scientific and engineering challenge, particularly in the development of clean combustion technologies based on hydrogen and other alternative fuels. Accurately predicting flame behaviour, pollutant formation, and combustion stability under practical conditions requires resolving a wide range of physical and chemical processes across multiple scales. The use of a spectral element method for direct numerical simulation (DNS) provides the necessary fidelity but is computationally demanding, especially for reacting flows with detailed chemistry. The deployment on extreme-scale computing systems would allow those computational needs to be met. However, this requires the solver to be optimised for execution on GPUs and in conjunction with the stabilisation of the numerical method, minimisation of communication, and highly efficient use of memory on the GPUs. The Ex3S (“Extreme Speed Scalar Solver”) innovation study focused on making solver improvements within the spectral element code NekRS. Performing DNS of hydrogen combustion at unprecedented scale and fidelity provides valuable data for model development and validation. This contributes to addressing the bigger challenge of delivering the predictive capabilities needed for designing next-generation, low-emission combustion systems.

Solution

Ex3S extended the GPU-accelerated spectral element solver, NekCRF, improving its ability to solve large sets of coupled scalar species with high efficiency and accuracy. By leveraging mixed-precision computing, advanced stabilisation techniques, and asynchronous communication, the solver overcomes key bottlenecks in performance and memory usage. The solver is capable of simulating turbulent hydrogen combustion with detailed chemistry at scale. Integrated into the NekRS framework via the NekCRF plugin, the solver demonstrated excellent performance, achieving near-roofline efficiency and strong scalability on leading European HPC systems, including best-of-class performance in the JUPITER early access call. Furthermore, the set of simulations performed resulted in a DNS database, which is among the first of its kind for hydrogen and serves as a reference for model validation.

Benefits & Impact

The new solver provides a powerful, scalable tool for simulating turbulent reacting flows with detailed chemistry. For the scientific community, it will enable fundamental research into flame dynamics, pollutant formation, and combustion instability. The generated DNS database for hydrogen combustion offers a unique benchmark for developing and validating next-generation combustion models. The availability of the new solver also has benefits for industrial users: it will support the design and optimisation of carbon-free propulsion and energy systems, particularly gas turbines and engines running on hydrogen or alternative fuels. Its ability to run high-fidelity simulations at realistic conditions reduces reliance on expensive experiments and accelerates development cycles. In general, the new technology can contribute to the achievement of climate goals, supporting clean energy strategies by improving predictive capabilities for emerging hydrogen technologies.

Next Steps

The Ex3S outcomes pave the way for future simulations of more complex fuels and chemical mechanisms, while establishing a robust foundation for advancing exascale combustion modelling in Europe.

Organisations Involved
Technical University of Darmstadt, Institute for Simulation of Reactive Thermo-Fluid Systems; Jülich Supercomputing Centre (JSC).

Contact: Prof. Dr.-Ing. Christian Hasse, Institute for Simulation of Reactive Thermo-Fluid Systems, hasse@stfs.tu-darmstadt.de

exaSIMPLE

Leveraging Machine Learning for Exascale CFD

Highlights

- Graph neural networks using GPUs were implemented as a replacement for conventional, computationally expensive solvers in the CFD algorithm SIMPLE.
- Using a data-driven approach improved performance and robustness in canonical benchmark and industry-typical CFD cases.
- The use of PETSc libraries and PyTorch makes the developments in this innovation study easily deployable in other CFD software. Open source code resulting from this innovation study is available in a GitHub repository for others to use.

Challenge

Initially developed in the 1970s, SIMPLE is a numerical algorithm that is included in widely used solvers for the Navier-Stokes equations in CFD simulations. Like many algorithms written for earlier HPC architectures, however, it is currently limited in its ability to take full advantage of the performance that current and future exascale systems offer. Solving the pressure correction equation in SIMPLE constitutes a major performance bottleneck, for example, accounting for between 60% and 90% of the computational time required to run a simulation due to its suboptimal use of memory bandwidth and communication patterns on today's massively parallel CPU and hybrid CPU-GPU hardware architectures. This limits scalability, ultimately constraining users' ability to run larger and more complex CFD simulations efficiently. Using machine learning (ML) algorithms to enhance the capabilities of CFD solvers is one strategy for addressing this problem that is currently attracting great interest. Data-driven models could replace or complement computationally expensive classical simulation methods, making the simulation of larger systems more tractable, reducing computational time and energy consumption, and avoiding the need to resort to simplified models or coarse meshes.

Solution

The innovation study exaSIMPLE introduced machine learning and data-driven models at two critical points within SIMPLE, focusing on simulation problems that are relevant for ocean, maritime, and renewable energy research and technology development. Firstly, the researchers developed a more efficient way to perform pressure correction in solving the Navier-Stokes equations, using graph neural networks to replace traditional solvers. Secondly, they made improvements in pressure-velocity coupling. This involved eliminating a key simplification made in SIMPLE while using ML to approximate a matrix inverse that is typically left out of SIMPLE simulations. Implementing this hybrid simulation approach on the EuroHPC JU supercomputer Deucalion, exaSIMPLE performed ML training on GPUs while running CFD simulation tasks on CPUs. All work was done using PETSc libraries in Python, a data structure for scalable scientific applications in MPI. This marked the first time that PETSc and machine learning methods were combined to push the boundaries of computational performance in CFD simulation.

Benefits & Impact

exaSIMPLE demonstrated that graph neural networks can be integrated into CFD solvers for solving linearised equation systems. The project derived a novel ML-based architecture and used it to solve typical linear pressure systems from benchmark, canonical, and industrial cases. This included calculating more than 7,000 test-cases to train the novel, graph neural network-based architecture. The trained algorithm infers solution vectors entirely on GPUs, simplifying use of hybrid CPU-GPU architectures for CFD. In addition, by reducing certain approximations commonly made in the traditional SIMPLE algorithm and approximating other terms that are typically omitted, this method reduces the number of iterations required to solve the Navier-Stokes equations. The study developed a new code that reproduces its results (https://github.com/blueOceanSustainableSolutions/exaSIMPLE). This will make it possible for other investigators to integrate this approach in their own CFD solvers.

Next Steps

exaSIMPLE shows clear promise, but further data and optimisation are needed to improve inference accuracy and performance. Its results could also offer opportunities in other scientific fields that rely on solving large matrix equation systems, including biomedical research, astronomy, and electromagnetics.

Organisations Involved
blueOASIS, INESC TEC, Marine Research Institute Netherlands (MARIN).

Contact: Guilherme Nuno Vasconcelos Beleza Vaz, blueOASIS, gvaz@blueoasis.pt

MG4ML

Maintaining Precision while Scaling Particle Physics Simulations

Highlights

- Multi-level, multigrid, and domain decomposition methods were combined to improve efficiency and accuracy in lattice quantum chromodynamics (QCD) simulations.
- Decomposition methods enable significantly higher scalability of the Dirac operator on GPUs in comparison to standard algorithms.
- Multigrid methods using coarse grids enhance the efficiency of low-mode averaging, a statistical noise reduction technique.

Challenge

Lattice quantum chromodynamics (QCD) is a numerical approach for studying interactions between quarks and gluons. Simulating QCD on supercomputers enables precise predictions of hadron masses, decay rates, and couplings, providing essential input for particle and nuclear physics experiments. As exascale architectures with heterogeneous components become available, lattice QCD algorithms must be optimised for efficiency and scalability. Lattice QCD simulations require high-resolution grids and extensive statistical ensembles. A critical step in these calculations, especially for studying hadrons, is inverting the Dirac operator to compute quark propagators. This operator is poorly conditioned, however, making solvers computationally expensive. At the same time, standard algorithms for this task face long convergence times and high noise levels, limiting precision. Over the past decade, research has explored the use of multigrid solvers as a way to address this challenge. Here, multiple levels of grid coarsening reduce problem size and accelerate convergence while preserving key information of the Dirac operator. This approach has been limited in its scalability, however.

Solution

The Inno4Scale study MG4ML set out to overcome this problem by integrating multilevel methods and novel multigrid-based algorithms within QUDA, a widely used library for accelerating lattice QCD calculations on GPUs. Multilevel block decomposition isolates active regions, enabling more frequent measurements while reducing statistical noise, leveraging locality, and preserving accuracy. The approach maintained strong scaling up to 192 GPUs, offering nearly double the scalability of standard approaches. The investigators also applied multilevel techniques to calculate correlations between two scalar loops, observing a variance reduction of up to a factor of eight. Additionally, multigrid solvers offered a more effective low-mode averaging technique when using coarse grids compared to direct eigensolvers. Together, these implementations enhance computational efficiency while lowering costs and are fully compatible with GPU-based exascale architectures.

Benefits & Impact

This study enhances the efficiency and precision of lattice QCD simulations, with potentially far-reaching benefits. Taking lattice QCD to exascale could help refine calculations in hadronic physics, improving theoretical predictions that can be compared with experimental data from particle accelerators at CERN and Jefferson Lab. In addition, the optimised domain decomposition and multigrid techniques developed in this study are applicable in fields beyond lattice QCD, including fluid dynamics, climate modelling, and materials science. The methods will help improve resource utilisation in national and international high-performance computing facilities, promoting more efficient and impactful computational research.

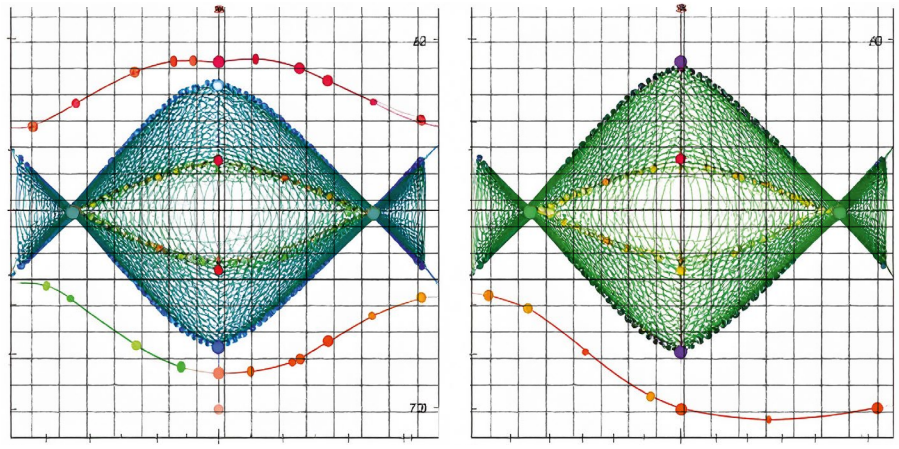

Next Steps

Additional tasks could enhance the overall performance and flexibility of this approach: Linear solver algorithms used to compute the inverse of the Dirac operator on a domain would profit from domain decomposition-aware reductions and level-1 BLAS routines. Support for even-odd reduced operators would also be beneficial, as these terms play a crucial role in the convergence of the solver. Performance enhancements could also be achieved by implementing iterative linear solvers in GPU kernels.

Organisations Involved
The Cyprus Institute, University of Wuppertal.

Contact: Simone Bacchio, The Cyprus Institute, s.bacchio@cyi.ac.cy

MLMC-Pint4Data

Taming uncertainty with multi-level Monte-Carlo and parallel-in-time methods

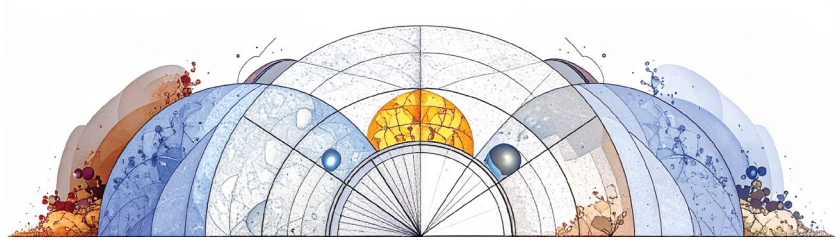

Highlights

- Multi-level Monte Carlo (MLMC) methods reduce computational cost of uncertainty quantification by combining simulations of varying accuracy.
- Integrating parallel-in-time (PinT) methods into MLMC reduces time-to-solution by enabling parallelism across higher numbers of compute cores.
- Combining MLMC and PinT methods will make uncertainty quantification more practical for large-scale simulations like those in fusion research and time-sensitive applications like weather prediction.

Challenge

Traditional simulations require exact values for all parameters necessary to define a problem. In practice, however, this often is not the case, as the parameter values obtained from measurements are subject to uncertainty. Uncertainty quantification (UQ) offers scientists methods to characterise and understand the effects of these uncertainties on the simulation result, but this also raises the computational effort considerably. This can become particularly challenging, for example, in applications in operational settings (e.g. weather prediction, computational medicine, or control optimisation), where time is of the essence. Here, it is typically possible to offer only simplified UQ or none at all. This challenge also arises in large-scale applications with many uncertainties (e.g. climate modelling, fusion energy research, or virtual prototyping), where UQ can be very expensive. The MLMC-PinT4Data innovation study investigated a new Monte Carlo (MC) method for uncertainty quantification. In Monte Carlo methods, the underlying problem is solved many times for random parameter values following a given statistical distribution. While MC methods are considered effective for problems with a high number of uncertain parameters, they are also computationally expensive (both in terms of energy use and time-to-solution). The goal of this study was to develop a Monte Carlo method that substantially reduces energy usage and time-to-solution compared to naive MC sampling.

Solution

In Monte Carlo sampling, the numerical solutions that result from individual parameter samples are independent of one another. This means that the time-to-solution can be reduced by using parallel computing to increase the number of samples that can be evaluated simultaneously. The total computational cost (energy-to-solution), however, remains unaffected. The Multi-Level Monte Carlo (MLMC) method employed in this project evaluates many samples using a cheaper, less accurate numerical method, retaining relatively few sample evaluations of the desired accuracy. In this way, the method preserves both the low bias error associated with accurate but costly sample evaluations and the low variance error that occurs when evaluating a high number of samples. Using such a mixture of sample evaluation levels also drastically reduces computational cost. While the MLMC method decreases the energy-to-solution, its overall parallelism and speed-up (time-to-solution) are still limited by the need to compute fine-level problems. The researchers increased the parallelism used to compute these problems by integrating an algorithm into their approach called Parareal, a so-called "parallel-in-time" (PinT) integration scheme. Unlike simulation methods that subdivide a domain into spatial units, Parareal subdivides a time interval under consideration into subintervals. Initial values for each subinterval are determined using a sequential but coarse and numerically cheap method, and the simulation for the subintervals is performed in parallel on different processing cores. Afterwards, the results of the individual time steps are integrated and a sequential correction is then used to reduce gaps at subinterval boundaries. This procedure is repeated iteratively until a certain tolerance is reached. In this way, Parareal reduces the overall time-to-solution of the Multi-Level Monte Carlo method by leveraging additional processor cores.

Benefits & Impact

The researchers applied the MLMC-Parareal solution to a Monte Carlo solver for the McKean-Vlasov equations, which are used to characterise stochastic processes in particle systems. The process they investigated was relevant for designing electrical machines and their solution is of high interest to leading manufacturers. To demonstrate the effectiveness of their approach with respect to uncertainty quantification, they showed that the MLMC-Parareal method reduces energy-to-solution to 8.3% of a standard Monte Carlo method. They also determined that their parallel-in-time method reduces time-to-solution to 57% of that of MLMC alone, with only a 1% increase in energy-to-solution (measured on a configuration with 8 nodes and 192 cores in total). As computing moves to exascale, these findings demonstrate the potential of parallel-in-time approaches to make uncertainty quantification more accessible for dynamic large-scale and time-critical applications of high-performance computing.

Next Steps

PinT methods offer an attractive speedup potential for problems for which a very long time-interval needs to be resolved, or a high resolution in time is required. In the context of finite element (FE) simulations, a high resolution in time is typically feasible only for simplified models in two spatial dimensions. For three-dimensional problems, the number of spatial Degrees of Freedom (DoF) can quickly become large enough to warrant additional parallelism in space, opening up another aspect of parallel computing. Further study of the optimal combination of sample-, space- and time-parallelism is required to enable (time-) efficient uncertainty quantification for three-dimensional problems.

Organisations Involved
Dept. Computer Science, NUMA, KU Leuven, CEM (Computational Electromagnetics) group, TU Darmstadt.

Contact: Giovanni Samaey, KU Leuven, giovanni.samaey@kuleuven.be

NEOSC

Cutting time-to-solution in half with an opportunistic data operations platform

Highlights

- Successful design, construction, and deployment of an opportunistic data operations platform (ODOP) enables on-the-fly data analysis using idling CPU resources in parallel with large-scale simulations, reducing the need for post-processing steps.
- Benchmarking tests of the ODOP pipeline on LUMI demonstrated a 50% reduction in time-to-solution in a turbulent convection simulation.
- ODOP has been deployed on LUMI. It offers domain scientists a tool for getting results more quickly and at lower computational costs, while enabling exascale HPC centers to increase energy savings and scientific throughput.

Challenge

While running large-scale fluid simulations using exascale supercomputers offers tremendous opportunities to achieve higher resolution over ever larger domains, it also raises a major challenge. Typically, data analysis takes place in a separate post-processing step after a fluid simulation is completed. At exascale, however, the enormous amount of data these simulations produce makes this impractical. This is because moving large-scale data sets over networks can be extremely time consuming, becoming a bottleneck that reduces performance while raising memory requirements and data storage costs. To improve time-to-solution, programmers of fluid simulations need new methods for efficient, on-the-fly data analysis and data management that can extract insights in parallel with simulation data generation.

Solution

One way to address this problem is to exploit the availability of idling CPU resources that are not actively being used in a simulation. In the innovation study NEOSC, researchers demonstrated the feasibility of a concept they call an opportunistic data operations platform (ODOP). This accelerated HPC pipeline is capable of executing fluid simulations at exascale, using Pencil Code interfaced with the Astaroth engine. During a simulation, ODOP monitors and schedules CPU resources not currently being used for data generation to conduct data analysis that in the past would have been done during post-processing. The framework includes a tool that enables domain scientists to identify optimisation opportunities with respect to execution time and computing costs without the need to make significant modifications to their code. NEOSC deployed ODOP on the EuroHPC JU system LUMI, and demonstrated its scalability and benefits in terms of time-to-solution in a real-world application related to turbulent convection.

Benefits & Impact

Benchmarking studies of ODOP showed that the pipeline can slash time-to-solution in a typical fluid dynamics simulation by half, in comparison to a workflow in which simulation and data analysis tasks were performed sequentially or in competition for the same node resources. Moreover, it was possible to complete data analysis tasks without any disturbance or slow-down of a GPU application on LUMI-G while using a large fraction of its GPU nodes (up to one-third of the full machine). For exascale system users, ODOP offers a solution that will help them to overcome the data analysis bottleneck and get results faster. For HPC centres operating exascale systems, this will also improve HPC resource utilisation, leading to energy savings as well as increased scientific productivity and throughput.

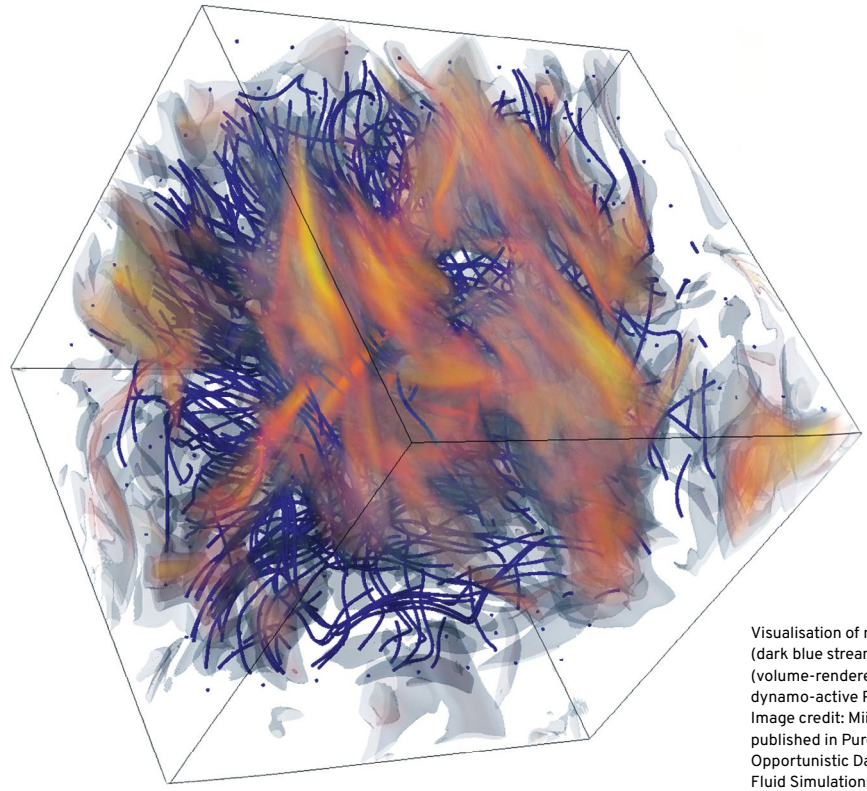

Visualisation of magnetic field lines (dark blue streamlines) and its intensity (volume-rendered colours) in a dynamo-active PC-A simulation. Image credit: Miikka Väisälä, to be published in Puro et al. 2025 "In-Situ and Opportunistic Data Analysis for Exascale Fluid Simulations", to be submitted to *Computer Physics Communications.*

Next Steps

The NEOSC tools are currently being integrated into CSC supercomputing facilities, and the first production runs are scheduled. The research team also submitted an application to the US Department of Energy's INCITE allocation program for access to computing time on the Frontier exascale supercomputer. The project would use tools developed in the NEOSC innovation study.

Organisations Involved
AaltoAstI & Aalto Systems and Services Engineering Analytics (AaltoSEA), Aalto University, CSC – IT Center for Science Ltd.

Contact: Maarit Korpi-Lagg, Aalto University, maarit.korpi-lagg@aalto.fi

ScalaMIDA

Improved Efficiency in Bayesian Inference for Earthquake Prediction

Highlights

- Prefetching and advanced model hierarchies improve scalability in uncertainty quantification.
- Hardware awareness optimises utilisation of CPUs and GPUs.
- Selective use of lower-resolution numerical simulations and data-driven surrogate models reduce the cost of large-scale inference.

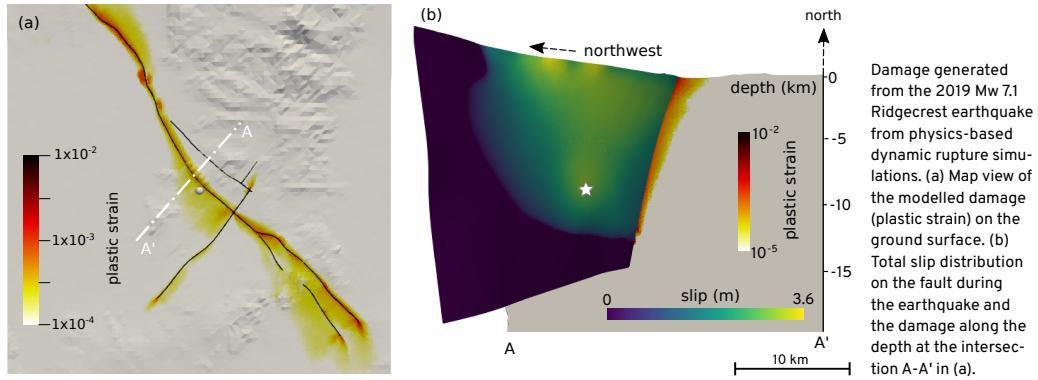

Damage generated from the 2019 Mw 7.1 Ridgecrest earthquake from physics-based dynamic rupture simulations. (a) Map view of the modelled damage (plastic strain) on the ground surface. (b) Total slip distribution on the fault during the earthquake and the damage along the depth at the intersection A-A' in (a).

Challenge

Bayesian inference is a statistical method for inferring otherwise inaccessible parameters from indirect observations, using a computational model and prior knowledge. In seismology, for example, seismographs on Earth's surface and satellite observations provide proxy data, but cannot by themselves deliver the comprehensive insights into complex behaviour at seismic faults needed to predict future earthquakes. Instead, seismologists use high-performance computing (HPC) to calculate what mathematicians call the Bayesian posterior; i.e., the distribution of likely parameters that define seismic activity based on available indirect observations. An important component of Bayesian inference is uncertainty quantification, which characterises the degree of confidence of the resulting predictions. This calculation is computationally expensive, because it requires many simulation runs, each requiring hundreds of compute nodes. This fact drives a need for greater efficiency and scalability both in statistical algorithms and HPC hardware.

Solution

In the innovation study ScalaMIDA, researchers tested methods for making Bayesian inference more efficient by retroactively simulating the complex fault system involved in the magnitude 7.1 Ridgecrest, California earthquake of 2019, based on measurements of ground motion and fault displacement. Algorithmic innovations included prefetching likely upcoming simulation runs in a Bayesian algorithm called Multilevel Delayed Acceptance (MLDA) and advancing the development of Multi-index Delayed Acceptance, an extension of MLDA, making use of higher-dimensional model hierarchies. In addition, they enhanced SeisSol, a software for simulating seismic wave phenomena, enabling efficient utilisation of CPU and GPU components through fused ensemble simulations. They also designed a hierarchy of approximate models for earthquake simulations in SeisSol, selectively using lower-resolution numerical simulations and data-driven surrogate models in place of computationally expensive simulation steps.

Benefits & Impact

In experiments performed within this study, prefetching MLDA enabled a more than 4× speed-up in the execution of small meshes on thousands of CPU nodes, while fused ensembles enabled up to 5× speed-ups in SeisSol. An ongoing focus on improving efficiency in Bayesian inference for large-scale simulation will make it possible to infer unknown parameters from data in complex, compute-heavy models, while accounting for model ambiguities and measurement errors. In practical terms, hardware-aware optimisations and hierarchical uncertainty quantification will improve SeisSol's performance on modern GPU-based supercomputers, enabling geophysicists to derive insights into seismic events at lower computational cost, improving forecasting and the ability to assist policymakers in risk assessment and mitigation. The novel approaches developed in this study are not just relevant for seismology, but could also lead to better insights in other fields that use Bayesian inference, such as climate science, engineering, and biomedicine.

Next Steps

Experiments in this study demonstrated that enabling prefetching is optimal in MLDA when scaling to large applications. Fused ensembles greatly improve the efficiency of low-order simulations that constitute the bulk of jobs for multi-level uncertainty quantification algorithms. Currently, limitations in fused simulations remain for high-order, large-scale models. Future work will focus on scaling these approaches to larger ensembles and simulation sizes.

Organisations Involved
Karlsruhe Institute of Technology, Technical University of Munich, Ludwig Maximilian University Munich.

Contact: Linus Seelinger, formerly Karlsruhe Institute of Technology, mail@linusseelinger.de

SCALE-TRACK

A New Particle Tracking Method for Cloud Simulations

Highlights

- A new highly parallel, efficient, asynchronous multi-GPU algorithm enables major scaling improvements in Lagrangian simulation in dispersed flow applications, with a 250× increase in the number of particles that can be simulated.
- Overall time to solution in combined Eulerian and Lagrangian simulations of clouds is improved by a factor of 4.5, while energy to solution is 3.4 times more efficient.
- Scaling runs on MareNostrum 5 reveal good weak and strong scaling.

Challenge

Understanding clouds is one of the big challenges of atmospheric physics and is important for climate science. Because investigating clouds in a controlled and repeatable manner is nearly impossible in the free atmosphere, researchers use cloud chambers, which make it possible to create and maintain artificial clouds in the laboratory. Numerical simulations also play important roles in cloud research, including in the construction of cloud chambers, the design and interpretation of experiments, and the examination of real-life clouds and climate phenomena. Cloud simulations – whether in nature or in cloud chambers – require a combination of two simulation regimes: The continuous flow of air is simulated using Eulerian methods in which a domain is subdivided into discrete units within a mesh. At the same time, the individual droplets or particles that make up clouds are simulated using a Lagrangian method. Here the movement of individual droplets is tracked as they move through the domain. Because the number of particles for cloud and climate simulations is very large, it is important that the simulation methods that investigators use scale well on modern HPC systems. However, in Lagrangian simulations it is currently very challenging to achieve a good scaling, primarily due to high communication overhead and non-uniform particle distributions, which lead to load imbalance across processors.

Solution

In this study, researchers improved the combination of Eulerian and Lagrangian methods for simulating dispersed flow phenomena by developing a new algorithm, called SCALE-TRACK, which overcomes the scaling issues connected to Lagrangian particle tracking simulation. Particles are processed in chunks, and the calculation of Eulerian and Lagrangian phases are executed asynchronously on CPUs (Eulerian) and GPUs (Lagrangian). This approach exploits the advantages of each architecture and reduces idle times. Furthermore, SCALE-TRACK uses a parallel decomposition in the Lagrangian phase that is independent from that of the Eulerian phase. In addition, an efficient implementation of the particles' properties limits memory requirements, which is particularly important for computing on GPUs.

Benefits & Impact

Before the start of the study, only one billion particles could be modelled with a simulation on a supercomputer, which is insufficient for realistic cloud simulations. SCALE-TRACK enabled the execution of calculations of that size on a local workstation. When deployed on the EuroHPC JU supercomputer MareNostrum 5, the new algorithm showed a remarkable performance boost, scaling efficiently up to 256 billion droplets. This advancement enables high-fidelity cloud chamber simulations, paving the way for much deeper understanding of these complex systems. The study also demonstrated that this method significantly reduces waiting times and synchronisation barriers, leading to a considerable speed-up. The overall time to solution is improved by a factor of 4.5, while energy to solution was 3.4 times more efficient than established methods. Scaling tests on MareNostrum (performed on up to 2,500 cores and 256 GPUs) revealed very good weak scaling (where the size of a problem per computing node is held constant) and strong scaling (where the total size of the problem is held constant). In addition to applications in cloud, climate, and weather simulation, SCALE-TRACK can be used for other instances of dispersed flow phenomena, where the dispersed phase consists of a large number of individual gaseous, liquid, or solid particles. This includes internal combustion engines, which spray a large number of fuel droplets during injection, or spray drying, as used in the chemical, pharmaceutical, and food industries.

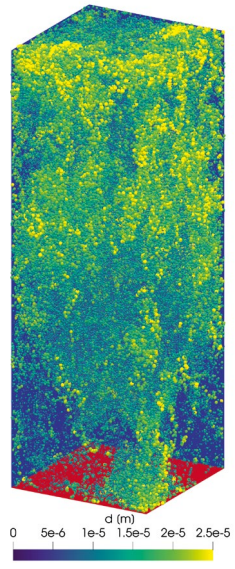

Snapshot of droplets in the 81m³ cloud chamber. Two side walls and the top are removed for illustration purposes and only a selection of droplets is shown. The size of the droplets is scaled up proportional to their diameter, but with an additional scaling factor to make them visible.

Next Steps

There are plans to build on the results already achieved by applying for a follow-up research grant. The goal is to further develop and expand the implemented algorithms and physical models, improving their applicability to real-world scenarios and thereby increasing their impact across various scientific communities, including weather and climate science as well as engineering.

TROPOS Leibniz Institute for Tropospheric Research

WIKKI

Organisations Involved
Leibniz-Institute for Tropospheric Research, Germany, Wikki GmbH Germany.

Contact: Dr. Henrik Rusche, Wikki GmbH, h.rusche@wikki-gmbh.de

XCALE

More powerful physics-inspired machine learning methods for materials research

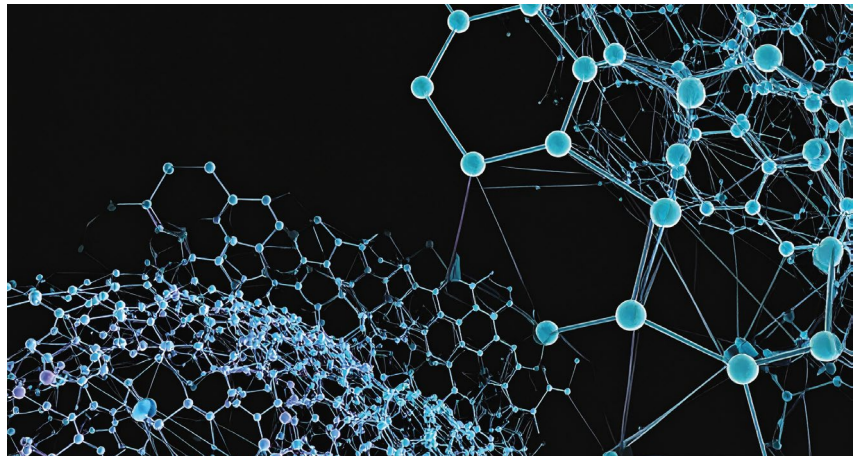

Highlights

- Machine learning offers a more computationally efficient approach over classic numerical methods for predicting atomic-level material properties. Integrating physics-inspired models into machine learning potentials improves accuracy, but algorithm development is needed to reduce computational costs.
- A new algorithm for many-body dispersion (MBD) theory and a new simulation paradigm for Generalised Hamiltonian Dynamics that offers 100× speed-up over existing methods will make it possible to simulate much larger molecular systems, increasing the feasibility of directly comparing simulation and experimental results in materials science.
- These new capabilities could support future research focused on improving efficiency in metal-ion batteries and catalysis in the production of green fuels.

Challenge

Traditionally, computational materials science runs complex simulations using numerical methods that accurately represent material properties down to the atomic level. These simulations are extremely computationally expensive, however. Machine learning potentials (MLP), which infer how atoms interact by predicting their potential energies based on their atomic structures, are emerging as an alternative approach, as they offer a fast and accurate way to explore how molecular interactions give rise to material properties. Purely data-based ML methods can be prone to inaccuracies, however, and so materials scientists have begun including physics-inspired interactions into MLP models to ensure that they adhere as closely as possible to physical laws. Unfortunately, this also comes at a significant computational cost. This new paradigm of atomistic modelling using MLP will become more practical and accessible by developing more efficient ways to add physics-based functionality, increase accuracy, and incorporate available experimental data for training. This multiscale, multimodal approach to atomistic modelling promises to bridge the gap between experiment and simulation in materials science.

Solution

Using TurboGAP, a parallel simulation engine for machine learning interatomic potentials on HPC systems, the innovation study XCALE achieved several advances in atomistic machine learning technology. These included implementing a new matrix-vector formalism for fast atomic-density descriptor calculations, developing ML-parameterised physical models for long-range electrostatic and Van der Waals interactions, and improving local-property models for Generalised Hamiltonian Dynamics, which match ML models to experimental X-ray diffraction data. This work also involved adapting and porting these methods to GPUs, an important step for taking MLP simulations to next-generation exascale systems.

Benefits & Impact

One key outcome of XCALE with high potential impact is a new algorithm that uses on-the-fly reparameterisation in an application of many-body dispersion (MBD) theory, an important approach for predicting the structure and stability of materials where weak bonding drives molecular structure or dynamics. Whereas previous methods rely on unrealistic pairwise approximations of dispersion, the new algorithm contains only minor approximations and improves scalability from $O(N^3)$ to $O(N)$ (that is, the time the algorithm requires to run is reduced from the cube of the number of atoms in the structure to just a linear relationship), allowing the accurate simulation of much larger molecular systems than were possible in the past without compromising accuracy. In addition, XCALE's implementation of a new simulation paradigm of generalised Hamiltonian dynamics for GPUs resulted in a ~100× speedup, with benchmarking done on up to 1,024 GPU cores. This could be transformative as it both integrates experimental data and relegates routines for characterising molecular systems dynamics to a minor computational expense, enabling experimentally valid atomic-scale structure prediction on much larger systems of molecules. Other improvements included achieving a 10× speed-up by implementing new radial expansion and kernel linearisation procedures, as well as a 5× speed-up in the implementation of electrostatics, including support for damped charges, using GPUs. Together, these achievements hold great potential to accelerate ML-based modelling of large-scale molecular systems, enabling more direct comparison of the results of atomic-level simulations and laboratory experiments.

Next Steps

Reflecting the potential applicability of the XCALE methodology for future work, the Aalto University group plans to study the solid electrolyte interphase, which can negatively affect the performance of metal-ion batteries. Using experimental measurements of the interface, they will derive atomic-scale models that will make it possible to diagnose the processes that cause problems such as the slowing down of ion diffusion. Other potential applications include the development of new catalysts for CO_2 reduction to green fuels or H_2 production from renewable energy sources.

Aalto University
School of Chemical Engineering

CSC

Organisations Involved
Aalto University School of Chemical Engineering, CSC – IT Center for Science.

Contact: Miguel Caro, Aalto University, miguel.caro@aalto.fi

Participants of the Innovation Studies and Inno4scale project consortium gathered during Innovation Day in Barcelona, May 2025, joined by EuroHPC Project Officer Linda Gesenhues.

The Inno4scale Consortium

Barcelona Supercomputing Center-Centro Nacional de Supercomputación (BSC-CNS) is the national supercomputing centre in Spain. It specializes in high performance computing (HPC) and manages MareNostrum, one of the most powerful supercomputers in Europe. BSC is at the service of the international scientific community and of industry that requires HPC resources.

scapos AG was founded in December 2008 for the sales, marketing and support of technical computing software. Technical software products from research institutions such as the Fraunhofer Gesellschaft and solutions from innovative start-ups characterize the scapos AG offering. scapos AG focuses on simulation (CAE), cutting and packaging optimization, metrology and language technologies. Its product portfolio includes: optimisation software, computer-aided engineering, libraries for HPC (in particular fast, scalable solvers for Linear Algebra problems). Its worldwide customers range from industrial corporations and large research institutes to SMEs. Its activities in HPC R&D projects complement the above and target service provision and the development of new business opportunities.

The High-Performance Computing Center Stuttgart (HLRS) was established in 1996 as Germany's first national high-performance computing center. As a research institution affiliated with the University of Stuttgart and a member of the Gauss Centre for Supercomputing, HLRS provides computing resources for academia and industry. HLRS operates HPC systems, provides advanced training, and conducts research. Among HLRS's areas of expertise are parallel programming, numerical methods for HPC, visualization, cloud computing, high-performance data analytics, and AI.

The Partnership for Advanced Computing in Europe (PRACE) is an international non-profit association (AISBL) with its seat in Brussels, Belgium. Developing as a Research Infrastructure since 2010, PRACE is currently implementing an organizational change process, shifting its core mission from providing access to Europe's largest supercomputers to representing the interests and needs of all HPC users in Europe. It is PRACE ambition to accelerate and expand the representation of user of HPC and related technologies (Artificial Intelligence, Quantum Computing, Cloud Computing, Data Science, etc) in Europe, and to enable and support high-impact scientific research and innovation developments across all scientific disciplines and industry, thereby enhancing European competitiveness.

The Inno4scale project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101118139. The JU receives support from the European Union's Horizon Europe Programme.

https://www.inno4scale.eu/
https://www.linkedin.com/company/inno4scale/
https://bsky.app/profile/inno4scale.bsky.social
https://www.youtube.com/@Inno4scale-xm7md/featured
https://www.inno4scale.eu/contact-us/

Design
GROOTHUIS. Gesellschaft der Ideen und Passionen mbH für Kommunikation und Medien, Marketing und Gestaltung; groothuis.de

Image Credits
p. 7, 11, 13, 15, 23: HLRS / generated with Firefly by Groothuis.
All other images: provided by the contact points of each study.
Cover page: unsplash / Aluminum Disemboweler3000