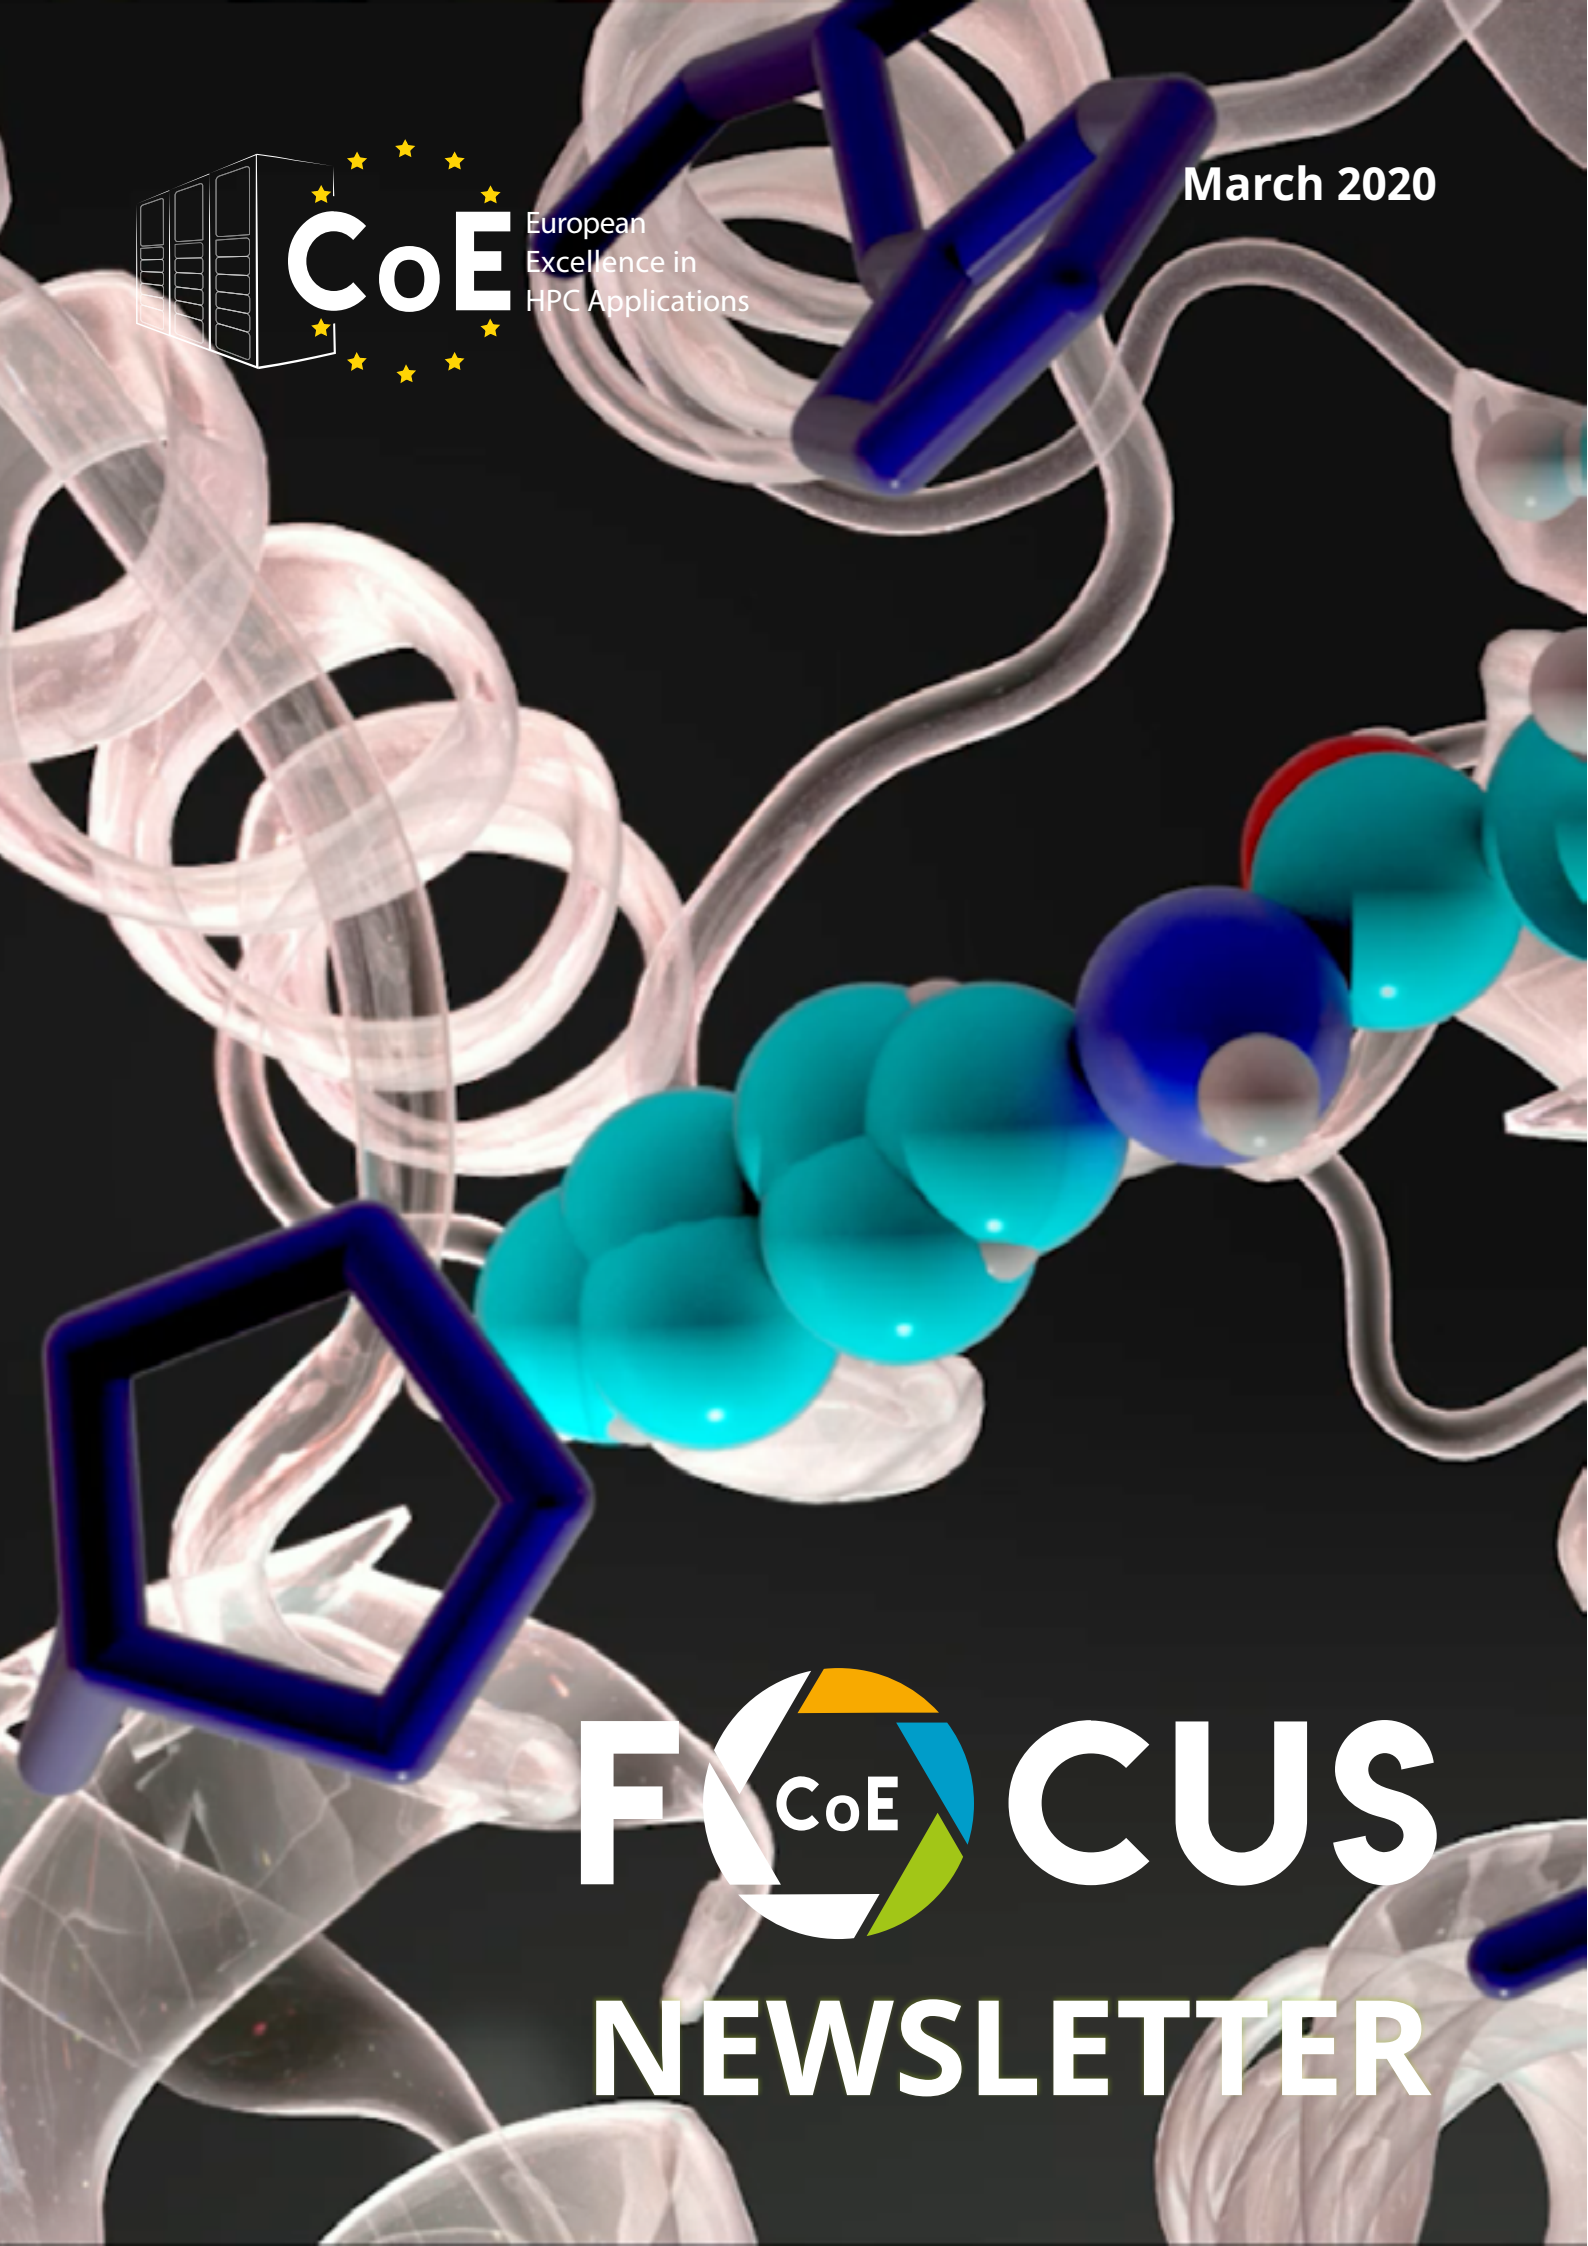
CoE
European
Excellence in
HPC Applications
March 2020
FOCUS
CoE
NEWSLETTER

WELCOME

Edouard Audit

chair HPC CoE Council

"The European HPC landscape is taking a major step forward with the creation of EuroHPC Joint Undertaking. It is allowing European countries to coordinate their effort and share resources in order to provide European engineers and researchers with a world-class supercomputing infrastructure, to foster the development of European technologies for HPC and to support the rich European applications ecosystem. As mentioned by Guy Lonsdale in the first FocusCoE newletter, HPC applications play a pivotal role; they transform aimless supercomputers into powerful tools that can address major industrial, scientific and societal challenges.

As we announced in the last newsletter, the active Centres Of Excellence (CoE) in high performance computing, supported by FocusCoE, established the HPC CoE Council – HPC3. One of the aims of HPC3 is to promote and ensure the strong support for application development within the HPC landscape shaped by EuroHPC. Applications are essential in order to yield science or innovation from computing resources and it's important that they receive long-term support.

In particular, HPC3 has been advocating for the implementation of a long-term vision of application development supported by long term funding. It should allow a closer collaboration between applications and hardware projects, especially the European Processor Initiative, in order to support a co-design approach; it should give proper support to large legacy codes that are key tools for many scientific and industrial communities and often need to be heavily refactored for new architectures.

European applications are great scientific assets that generate benefits for all areas of research and industry, they should be supported to engage in the best possible way in the Exascale era. Therefore, in the coming months in which the next European research agenda will be shaped and with the support of all CoEs, HPC3 will continue to advocate for a strong, ambitious, long-term and coherent policy for HPC applications in Europe."

CONTENTS

SUCCESS STORY: Drug Discovery

CompBioMed, Centre of Excellence in computational bio medicine, worked together with Janssen Pharmacentica NV to improve the molecular simulations in drug discovery.

1. Organisations involved:

SURFsara is the National Supercomputing and e-Science Support Center in the Netherlands. SURFsara provides expertise and services in the areas of High Performance Computing, e-Science & Cloud Services, Data Services, Network support, and Visualisation.

Janssen Pharmaceutica NV is an affiliate of the Pharmaceutical branch of the US Johnson & Johnson company. Janssen's interests are in developing and using advanced molecular simulation methods to optimise lead compounds in discovery programs, predicting the activity of compounds with specific targets.

2. Challenge:

Janssen's primary challenge is in developing and using advanced molecular simulation methods to optimise lead compounds in discovery programs. Such methods, if proven robust and accurate could have a profound impact on the way drug discovery is performed. They would permit reliable computational triaging of very close analogue molecules greatly improving efficiency. Also, this would lead to high-confidence design of synthetically more challenging molecules leading to better drugs in new chemical space.

3. Solution:

Through the CompBioMed HPC allocations service JAN obtained an allocation of 1.8M core/hours on Cartesius (SURFsara). JAN evaluated the use of non-commercial software for predicting the free energy of a series of compounds and built a protocol which uses open-source software for MD and free energy perturbation simulations. The main simulation code used was GROMACS (version 2016.1) configured to run on GPU, and the results compared to previous simulations performed with Schrodinger's FEP+ commercial software. Slurm support for job array and job dependencies, has been used to orchestrate and streamline the workflow and to optimise the MD simulations execution on multiple compute nodes.

4. Business impact for the private company:

New computational approaches described above can have a major impact on the drug discovery process. Importantly the work studied here will go to the heart of the design-make-test cycle and contribute higher quality methods for compound prioritisation.

This is a fundamental issue of drug discovery, because, whilst compounds with acceptable potency can often be found, and found quickly, they do not always come with other desired properties. Hence, during a typical lead optimisation (LO) program the challenge often becomes to maintain potency whilst modifying the chemical structure of the lead molecules to overcome these other issues. In this regard computational tools which can accurately predict binding mode combined with accurate binding affinity prediction will be extremely powerful and reduce the number of 'backwards' steps required to subsequently move forwards in an LO program.

The project has introduced Janssen to numerous research groups throughout Europe that are able to assist them in their aims and to combine expertise to produce useful results. The use of HPC is critical for these experiments, without which the compute power does not satisfy the needs of such complex systems.

5. Benefits:

- Higher quality methods for compound prioritization.
- Increased efficiency during lead optimisation program.
- Reduction in time/cost of synthesising redundant drug molecules
- Higher probability of determining new active drug

CompBioMed is a user-driven Centre of Excellence in Computational Biomedicine. They have users within academia, industry and clinical environments and are working to train more people in the use of their products and methods.

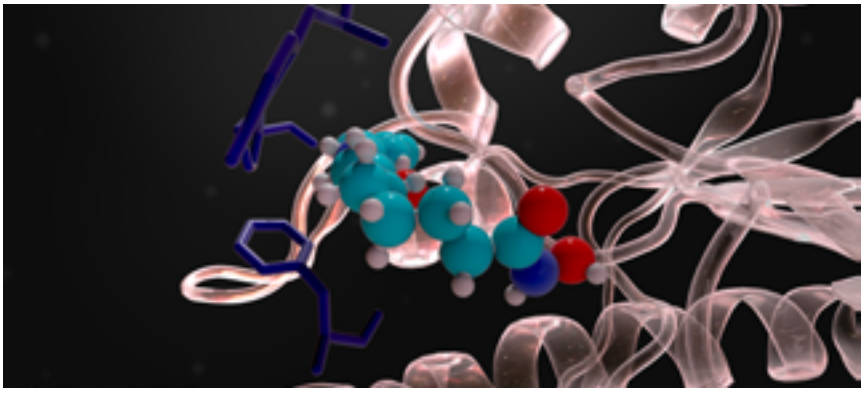

SUCCESS STORY:

Designing control pulses for superconducting qubit systems

Together with IBM research, the European Centre of Excellence E-CAM developed a software for designing control pulses to manipulate qubit systems.

1. Organisations involved:

Centre Européen de Calcul Atomique et Moléculaire CECAM, located at the EPFL in Lausanne, is an organisation devoted to the promotion of fundamental research on advanced computational methods and to their application to important problems in frontier areas of science and technology.

IBM Research Laboratory - Zurich is the European branch of IBM research, which is the research and development division of the American multinational information technology company IBM.

2. Challenge:

The aim of this pilot project was to develop a new method and dedicated software for designing control pulses to manipulate qubit systems (see Fig.1A) based on the local control theory (LCT) The system is composed of two fixed frequency superconducting transmon qubits (Q1 and Q2) coupled to a tunable qubit (TQ) whose frequency is controlled by an external magnetic field. Changing the frequency, the TQ behaves as a targeted quantum logic gate, effectively enabling an operation on the qubit states. The system schematises an approach to construct real quantum universal gates currently investigated by IBM.

3. Solution:

Local control theory (LCT), the main theoretical tool used, originates from physical chemistry where it is used to steer chemical reactions towards predetermined products, but it had never been used to design a quantum gate. To create the software, researchers added new functionalities to the open source QuTip program package. Two main modules were developed during the project: LocConQubit, which implements the LCT and accompanying procedures, and OpenQubit, a patch to the first module which introduces Lindblad master equation propagation scheme into the LCT which also enables direct construction of pulses under the presence of decoherence effects. All modules were written in Python and expand the functionalities of the QuTip programme package.

4. Business impact for the private company:

The developed software was transferred to IBM and will be of use to engineer pulses for the experimental construction of superconducting Qubits. A paper in collaboration with the industrial partner was recently published „Local control theory for superconducting qubits“ (M. Mališ, P. Kl. Barkoutsos, M. Ganzhorn, S. Filipp, D. J. Egger, S. Bonella, and I. Tavernelli, Phys. Rev. A 99, 052316).

5. Benefits:

• This project introduced a new procedure for generating control pulses on-the-fly in qubit systems, which is less computationally demanding and might thus open new approaches to pulse constructions

• LCT pulse can also serve as initial guess pulses in optimal control theory

• With LCT it was possible to design extremely short pulses with a duration of just a few tens of nanoseconds or less with an almost full fidelity (see Figure 1B).

• The method is highly robust and requires only qubit parameters (frequencies, coupling terms) as inputs

The European HPC Centre of Excellence (E-CAM) is an e-infrastructure for software development, training, and industrial discussion in simulation and modelling that started in October 2015. E-CAM focuses on four scientific areas of interest to computational scientists: Classical Molecular Dynamics, Electronic Structure, Quantum Dynamics, Meso- and Multi-Scale Modelling.

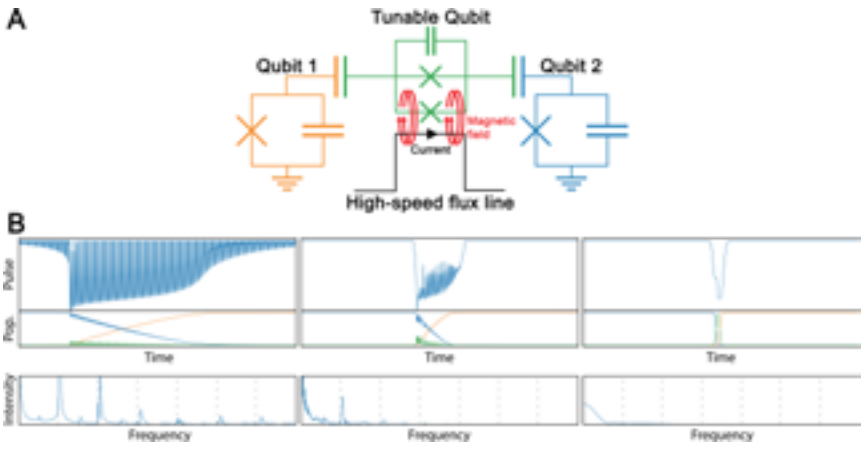

Figure 1: A) Schematic representation of the qubit system; B) Pulses obtained by LCT (left), by a frequency filtering LCT procedure (middle), and by analytic function fitted to LCT pulse parameters (right) with their corresponding population transferring from qubit 2 (blue) to qubit 1 (orange), and frequency spectra. FocusCOE contributes to the success of the EU HPC Ecosystem and the EuroHPC Initiative by supporting the EU HPC CoEs to more effectively fulfil their role within the ecosystem and initiative: ensuring that extreme scale applications result in tangible benefits for addressing scientific, industrial or societal challenges

SUCCESS STORY:

Harness the power of vorticity for next-gen wind turbines

The European Centre of Excellence E-oCoE teamed up with Spanish tech Startup Vortex to improve their next generation wind turbines.

1. Organisations involved:

VORTEX BLADELESS is a Spanish tech startup that is developing an environmentally friendly aerogenerator which needs no blades. It is a new wind energy technology especially designed for on-site generation in residential areas, being able to work on-grid, off- grid, or along with regular solar panels or other generators. The *Barcelona Supercomputing Centre* is a Spanish public research center that developed the multi-physics simulation code ALYA and provided the MareNostrum supercomputer.

2. Challenge:

Vortex Bladeless aims to harness the power of vorticity for a new generation of wind turbine. They develop a single column without bearings or gears. It just oscillates with the wind. Experiments with scaled down prototypes have been encouraging, but the physics behind these devices is highly complex. There is a need to optimise and explore scalability due to the complexity of the flow and the need for time accurate results. The needed Large Eddy Simulation (LES) simulations are computationally demanding.

3. Solution:

The company has been working with experts at the Barcelona Supercomputing Centre on the MareNostrum supercomputer. The fluid-structure interaction (FSI) between the Vortex Bladeless device and a turbulent flow is simulated with Alya.

The results from initial simulations of a scaled-down device were very close to the actual wind tunnel tests performed by the Vortex Bladeless team, allowing them to develop the idea of a range of devices at the micro scale and the utility scale. Then, the behaviour of the device at a more realistic scale was studied by means of numerical simulations, helping in the design of real scale experiments and reducing costs.

4. Business impact for the private company:

In order to understand the aerodynamic behaviour of the devise, experimental, or numerical studies can be performed. For the real scale model, that is approximately three meters high, experimental studies would have been too expensive. It was therefore decided to resort to high accuracy numerical simulations of the flow coupled to the oscillating turbine. Since the flow is highly transient, and it is important to capture the dynamically important scales of the flow accurately, Large Eddy Simulation (LES) was selected as the modeling technique.

Modeling of turbulent flow is still nowadays one of the most computational demanding problems. The interaction with the Barcelona Supercomputing Centre was, therefore, crucial to develop a much better understanding of the aerodynamics of the device. BSC provided the computation resources and the inhouse code Alya that can efficiently run on high-end Supercomputers. Moreover, they contributed with their know-how on LES turbulence modeling and Fluid-Structure Interaction. Technological advantages the HPC provides are important for small to medium-sized enterprises (SME) to remain competitive. For instance, the savings regarding the CPU time are key for SMEs to be able to use highly advanced techniques such as Large Eddy Simulation.

5. Benefits:

- Reduced costs by preparing the real scale experiments by means of numerical simulations
- The highly optimised code enabled to avoid throwing away costly computational resources
- The cpu time for assembly has been reduced up to 38%.
- A new solver has provided speed ups of up to five times with respect to Alya's own solvers.

EoCoE (Energy Oriented Center of Excellence: toward exascale for energy) is an energy-oriented Centre of Excellence for computing applications that builds on its unique expertise at the crossroads of high-performance computing (HPC) and renewable energy. It brings an impulse to accelerate the digitization of the future energy system. The coding developments are assisted by multi-disciplinary teams with expertise in applied mathematics and high-performance computing (HPC)

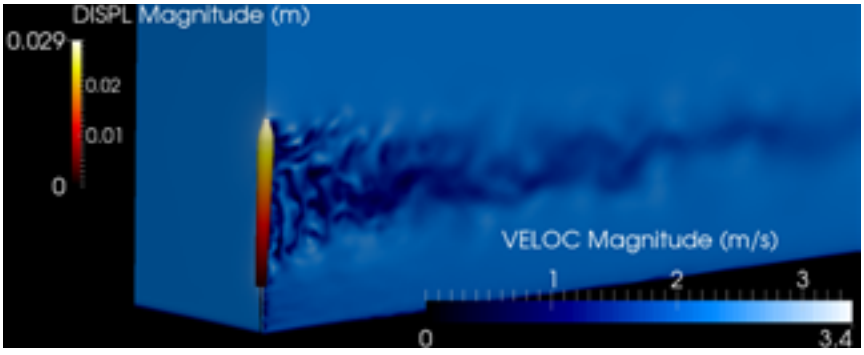

Simulation of the fluid-structure interaction (FSI) between the Vortex Bladeless device and a turbulent flow from Alya

NEWS:

Cluster Review and HPC3 meeting in Luxemburg

On December 17th and 18th 2019, representatives of the ten European Centres of Excellence in high-performance computing (CoE) took part in a cluster review at the European Commission in Luxemburg.
In addition, Focus CoE joined the meeting as the corresponding "Coordination and Support Action" for the CoEs. Besides the actual presentations on each of the project's progress and achievements, the participants discussed topics like HPC training and education, scalability and exascale computing in a world café format. The review is part of the regular reporting measures that the CoEs conduct throughout the runtime of a EU-funded project.
The CoE representatives took the opportunity in Luxemburg to also gather for a meeting of the HPC CoE Council (HPC3) that was founded last year.
Besides the preparation for the cluster review, the council set the course for a joint working group on business development and discussed measures around the CoEs training activities.

CODES:

ChEESE

SPECFEM3D
solves linear seismic wave propagation (elastic, viscoelastic, poroelastic, fluid-solid) and dynamic rupture problems in heterogeneous 3D models. SPECFEM3D also implements imaging and FWI for such complex models based on an L-BFGS (Broyden-Fletcher-Goldfarb-Shanno) algorithm. Based on the high-order spectral-element (CG) discretization for unstructured hexahedral meshes. Scalable performance at Petascale (runs on the largest machines worldwide: Titan and Summit at Oak Ridge, Piz Daint, CURIE, K computer, etc.). ChEESE has recently optimized SPECFEM3D which resulted in a two-fold speedup of the code.

EoCoE

Alya
is a high performance computational mechanics code that solves complex coupled multi-physics problems, mostly coming from the engineering realm. Alya will be deeply refactored in order to be able to address heterogeneous computing nodes. A very strong effort on linear algebra is planned for both the fluid and the solid problems handled in the two tasks foreseen in the wind scientific challenge.

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Nemo
is a modelling framework for oceanographic research, operational oceanography seasonal forecast and climate studies. It has three major components: NEMO-OCE models the ocean (thermo)dynamics and solves the primitive equations, NEMO-ICE models sea-ice (therms)dynamics, brine inclusions and subgrid-scale thickness variations, NEMO-TOP models the (online and offline oceanic tracers transport and biogeochemical processes. Several built-in configurations are provided to evaluate the skills and performances of the model. They can be used as a first easy set-up and as template for setting up a new configuration. Idealized test cases addressing specific physical processes are also available for download.

COEs OVERVIEW:

BioExcel 2 - Biomolecular Research
BioExcel is operating towards advancement and support of the HPC software ecosystem in the life science domain. Research and expertise covers structural and functional studies of the main building blocks of living organisms (proteins, DNA, membranes etc.) and techniques for modelling their interactions ranging from quantum to coarse-grained models up to the level of a single cell.

ChEESE - Solid Earth
In CheEESE, leading European HPC centers, academia, hardware developers, as well as SMEs, industry and public governance bodies such as civil protection are working together to prepare European flagship codes for upcoming pre-Exascale and Exascale supercomputing systems to tackle global challenges in the domain of solid earth.

CompBioMed - Biomedicine
CompBioMed is a user-driven Centre of Excellence in Computational Biomedicine, to nurture and promote the uptake and exploitation of high performance computing within the biomedical modelling community, supporting users in academia, industry and clinical practice.

E-CAM - Method and algorithm deployment
The overall objective of E-CAM is to create, develop and sustain a European infrastructure for computational science applied to simulation and modelling of materials and of biological processes of industrial and societal interest.

EoCoE-II - Energy
EoCoE will use the prodigious potential offered by the ever-growing computing infrastructure to foster and accelerate the European transition to a reliable and low carbon energy supply via targeted support to four carbon-free energy pillars: Meteorology, Materials, Water and Fusion, each with a heavy reliance on numerical modeling.

ESiWACE2 - Weather & Climate
ESiWACE2 aims to link, organise and enhance Europe's excellence in weather and climate modelling to enable leading European weather and climate models to leverage the performance of pre-exascale systems as soon as possible and prepare the weather and climate community to be able to make use of exascale systems when they become available.

EXCELLERAT - Engineering
EXCELLERAT's goal is to facilitate the development of important codes for high-tech engineering, including maximizing their scalability to ever larger computing architectures and supporting the technology transfer that will enable their uptake within the industrial environment.

HiDALGO - Global Challenges

HiDALGO enables the assessment of Global Challenges problem statements by enabling highly accurate simulations, data analytics, artificial intelligence and data visualisation, but also by providing knowledge on how to integrate the various workflows and the corresponding data.

MaX - Materials Design

MaX aims to disenthrall the EU leadership in materials modelling, simulations, discovery and design by creating an ecosystem of capabilities, software applications and data workflows and analysis on HPC-oriented material simulations, designed for present and future HPC architectures.

POP-2 - Performance optimization & productivity

POP-CoE assesses the performance of computing applications, identifying issues affecting code performance as well as the best ways to address them.

Funded under the Horizon 2020 research & innovation programme from European Union

www.focus-coe.eu

Twitter: @FocusCoE

FOCUS CoE

European Excellence in HPC Applications