EXCELLERAT

NEWSLETTER

2020-05

Editor's note

What is EXCELLERAT?
EXCELLERAT is an EU project that combines key players from industry, academia and High Performance Computing (HPC) centres with the necessary European expertise to establish a Centre of Excellence in engineering.
EXCELLERAT's mission is to pave the way for the evolution towards Exascale technologies. Exascale computing will have the power to address highly complex and costly engineering problems, creating insight into the potential of technological solution. We provide service solutions for industry and research, providing the knowledge, computational power, and infrastructure crucial to tackle the ever-rising complexity of scientific and development endeavours. Our Service Portal offers a range of services such as trainings and individual consultancy that will help companies or researchers incorporate HPC into research or product development. EXCELLERAT will provide users with tailored support they will benefit from individually, based on their own current needs.

In this first newsletter issue, we would like to share stories about our partners' support in the fight against COVID-19 and a wrap-up on the HPC User Forum, introduce some of the partners we collaborate with and give our readers an overview of EXCELLERAT news and upcoming events. Don't forget to subscribe to the Newsletter and enjoy reading!

Your EXCELLERAT team

Contents

Our partners support the fight against COVID-19

The number of persons infected by the COVID-19 virus is increasing globally at a remarkable rate; in many places hospitals are running at capacity and facing the threat of shortages of essential supplies. The impact of the new coronavirus SARS-CoV-2 reminds us of science fiction scenarios which until recently were hard to imagine in real life. In order to contain the coronavirus' effects, several of our project partners – including CINECA, the Barcelona Supercomputing Centre (BSC), the Edinburgh Parallel Computing Centre (EPCC), RWTH Aachen University and the High-Performance Computing Center Stuttgart (HLRS) – are supporting measures to fight the pandemic.

As part of the EU-funded project Exscalate4CoV, our partner CINECA, together with BSC, KTH and the Forschungszentrum Jülich, are helping analyse COVID-19 proteins based on data from the scientific community in order to accelerate txhe search of an effective therapy. Carlo Cavazzoni, senior staff member at CINECA said in *an interview with This Week in HPC*, "All the research teams working around these supercomputing centres – the researchers and the scientists – are sharing knowledge, working together, to ensure that we get the fastest possible good candidate drug for the coronavirus."
BSC is collaborating in the fight against the coronavirus using several different approaches: applying bioinformatics to better understand the virus and identify possible treatments, using Artificial Intelligence (AI) and natural language processing to analyse the data about the spread and impact of the pandemic, and making the MareNostrum 4 supercomputer available for the fight against the coronavirus. Bioinformatics approaches involve analysing the coronavirus genome and its successive mutations, and searching for drugs and immune therapies (antibodies and vaccines) that could most effectively target key vulnerabilities in the virus. BSC's High Performance Artificial Intelligence (HPAI) research group is collaborating with UNICEF and IBM on a project to analyse the socioeconomic impact of the virus locally and globally, with an emphasis on social distancing. The goal is to find impact indicators, patterns and statistics that could help the UN and local authorities to take better and faster measurements.
The MareNostrum 4 supercomputer is also providing researchers at BSC and in the wider scientific community with the necessary computational capacity to accelerate ongoing investigations against the coronavirus. *EPCC* is part of the Rapid Assistance in Modelling the Pandemic (RAMP) initiative, which supports existing COVID19 work by enhancing the modelling teams who inform UK Government policy. *The Fab Lab at RWTH Aachen University* has made its 3D printers, laser cutters, and other digital production machines available for the rapid production of urgently needed parts like visors for doctors, emergency ventilators for hospitals, and forearm door handles for office doors. These are currently being prototyped as part of a cooperation between the Fab Lab and the Uniklinik RWTH Aachen.
HLRS and its partners in the Gauss Centre for Supercomputing have announced that they will expedite access to HPC resources for scientists pursuing research aimed at prevention, containment, remediation, or cures related to the coronavirus pandemic. *HLRS has committed to fast-tracking applications for COVID-19 related computing time, minimising all hurdles during the application process*. This applies to research at the molecular level to understand the virus and develop vaccines and therapeutics, epidemiological research to understand and forecast disease spread, and other related approaches aimed at halting the pandemic.
Apart from these examples, other *European supercomputing centres and initiatives* have also offered their support during the COVID-19 outbreak. Let's hope that these efforts will lead to solutions and an improvement of the current situation.

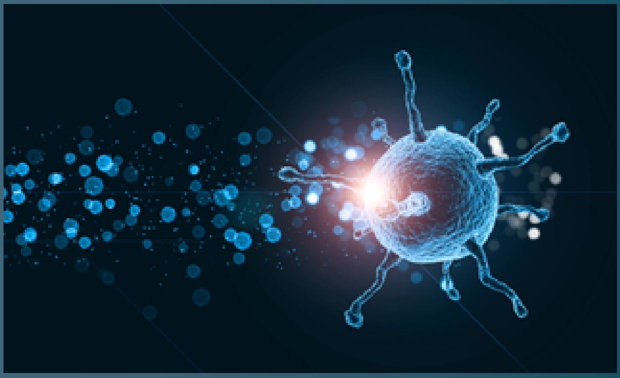

Report on iHurt Industrial HPC User Roundtable

In addition to enabling academic research, the High-Performance Computing Center Stuttgart (HLRS) supports industry by making its supercomputing resources for simulation, machine learning, and visualization available for research and development. To better understand and address the specific needs of industrial HPC users, SICOS BW and HLRS hosted the third annual Industrial HPC User Roundtable (iHURT) on December 3, 2019. The event offered a forum for HLRS and its industrial users to exchange perspectives on the state of the art in high-performance computing, specific computing challenges industrial users face, and how HLRS could help address them.

The program opened with presentations by industrial users of HLRS's HPC systems. Andreas Link of automotive supplier Mann + Hummel described his team's computational fluid dynamics (CFD) simulations, which they run at HLRS to optimise car component geometries and reduce noise caused by air flow around them. When combined with rapid prototyping technologies, the approach makes it possible to test new part geometries quickly. René Thümmler of CFD Consultants GmbH, a small consulting firm that assists companies in running CFD simulations, talked about his experience using HLRS computing resources, pointing to challenges the company faces with respect to data transfer and software licensing.

This year's iHURT meeting paid particular attention to growing interest in the integration of artificial intelligence applications and high-performance computing. Iris Pantle of a startup called Falquez, Pantle und Printz GbR described her company's goal of integrating machine learning and classical CFD simulation. The company is in the early stages of developing a cloud-based platform that would use machine learning to identify key parameters in CFD data sets, analyse simulation results, and optimise product development in an interactive, automated way. Alexander Thieß of Spicetech GmbH described how his company provides AI as a service for other companies interested in using machine learning to extract insights from the data they collect.

Representatives of HLRS and SICOS BW described recent developments at HLRS that are relevant to industrial HPC users. SICOS BW Executive Director Andreas Wierse provided an update on EXCELLERAT. Thomas Bönisch, who leads the HLRS user support team, introduced the HPC center's new flagship supercomputer, called Hawk, which is due to go online early in 2020 and will offer more than triple the computing power of its outgoing system, Hazel Hen. SICOS BW's Nicole Dobner provided an update on the Supercomputing-Akademie, a continuing professional education program involving HLRS, Ulm University, and University of Freiburg that is designed to address the specific needs of industrial HPC users.

Capping the meeting was an open discussion focusing on practical challenges that industrial HPC users face. Topics discussed included training, data management, workflow development and data transfer, access to software licenses, and the integration of HPC and artificial intelligence.

Partner: EPCC

EPCC is the high-performance computing (HPC) centre at the University of Edinburgh. Founded 30 years ago, it has grown to be the UK's largest HPC centre and one of the biggest in Europe.

EPCC has multiple roles in EXCELLERAT. We are developing the application code TPLS (which models the two-phase flows often encountered in chemical engineering processes) with the aim of achieving the high scalability that will be needed for Exascale. We are also providing expertise on advanced meshing techniques and the use of novel architectures such as FPGAs, which will be a key element of Exascale systems.

EPCC works closely with industry. Recent collaborations include the Fortissimo project in which over 100 companies (mainly SMEs) were supported in running HPC trials. We currently co-lead the UK's ASiMOV consortium with Rolls-Royce. This five-year programme was awarded £14.7 million to develop the next generation of engineering simulation and modelling techniques, with the aim of developing the world's first high-fidelity simulation of a complete gas-turbine engine during operation.

We are at the forefront of the European drive towards Exascale. Current projects include VESTEC, which is investigating new models for "urgent supercomputing" and real-time data feeds with the aim of using supercomputers as decision-support tools for natural-hazard emergencies such as wildfires. SAGE2 is investigating how to manage the enormous volumes of data created by advanced computing techniques. EPiGRAM-HS is improving and extending existing programming models with Exascale potential, with the ultimate aim of making complex machines easier to use efficiently.

EPCC hosts and manages an exceptional range of advanced computing systems. ARCHER, the UK's supercomputing service for research, and the UK Research Data Facility are both located at EPCC's Advanced Computing Facility. The £79 million follow-on ARCHER2 service will also be hosted by EPCC. Other HPC systems include Cirrus, which provides an ideal platform for solving computational simulation and modelling challenges. Both ARCHER and Cirrus are available for industry use on a pay-as-you-go basis.

Finally, we are currently building the Edinburgh International Data Facility (EIDF), which will offer services for the long-term hosting and preservation of digital data and for state-of-the-art analytics and data science. This major facility will host a rich and growing collection of data and will also provide safe haven services to health and government users.

EPCC: *www.epcc.ed.ac.uk*
EIDF: *www.ed.ac.uk/edinburgh-international-data-facility*

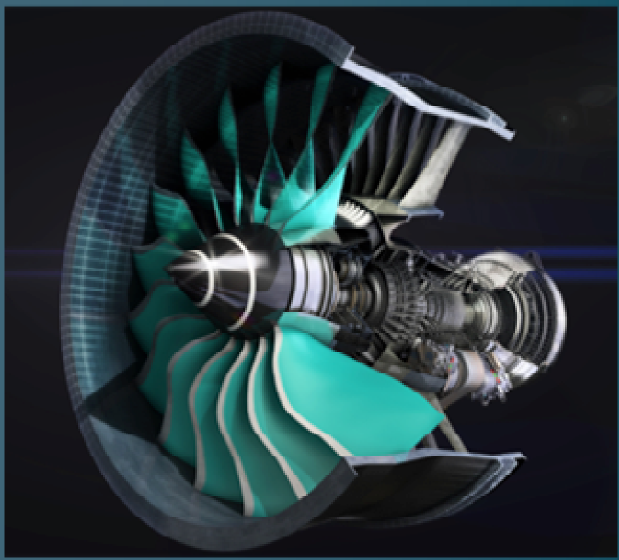

Partner: SSC-Services GmbH

Creating Connections - We connect working worlds.

As an IT service provider, SSC develops individual concepts for cooperation between companies and customer-oriented solutions for all aspects of digital transformation. Since 1998, the company, based in Böblingen (Germany), has been offering solutions for the technical connection and cooperation of large companies and their suppliers or development partners, mainly in the automotive industry.

Our corporate roots lie in the management of data of all types and sizes. SWAN – our own developed and operated data exchange platform – is a modern, secure and reliable solution, which has been successfully in use for more than 20 years and is constantly being further developed. As an industrial partner in the EU project EXCELLERAT, we are also promoting the topics of data management and data transfer. Part of EXCELLERAT's vision is to provide the engineering community with easy access to relevant services and knowledge in the field of high-performance computing.
We are developing the platform for scientific and academic HPC users and medium-sized industrial companies. The platform is not only used for data processing, but also enables bidirectional, online data transfer between the data generators and all six high-performance computing centres represented in the EU project. This data transfer will be highly automated in order to avoid duplication of the transferred content. This approach will lead to a data reduction, which can ultimately save time and costs.
SSC currently provides a secure workplace for 170 employees from 21 countries of origin. We strive for a diversity of our employees, customers and projects. In order to be able to sustainably succeed in the global market, we welcome the change. Change requires innovation. With agile methods, DevOps principles, open source technologies and software as well as automation processes, we use degrees of freedom for technical development.

Our goal is to eliminate manual effort from previous processes and to use human labour and creativity where it is needed. We like to break up old habits, question them and look for hidden potentials to improve and develop them.

Vision: SSC enables the bidirectional data exchange combined with intelligent data dispatching between scientific and academic HPC users / SMEs and European HPC Centres.

Dynamic Load balancing for Airplane simulations on Heterogeneous pre-Exascale architectures

We can affirm with quite a certainty that future Exascale systems will be heterogeneous, including accelerators such as GPUs. We can also expect higher variability on the performance of the various computing devices engaged in a simulation; due to the explosion of the parallelism, or other technical aspects such as the hardware-enforced mechanisms to preserve the thermal design limits. In this context, dynamic load balancing (DLB) becomes a must for the parallel efficiency of any simulation code.
In the first year of the EXCELLERAT project, Alya has been provisioned with a distributed memory DLB mechanism, complementary to the node-level parallel performance strategy already in place. The kernel parts of the method are an efficient in-house SFC-based mesh practitioner and a redistribution module to migrate the simulation between two different partitions online. Those are used to correct the partition according to runtime measurements.
We have focused on maximising the parallel performance of the mesh partition process to minimise the overhead of the load balancing. Our approach was presented in the SC19 conference in Denver [1]. We demonstrated that our software can partition a 250M elements mesh of an Airplane with 0.08 sec using 128 nodes (6144 CPU-cores) of the MareNostrum V supercomputer – ten times faster compared to the Zoltan software from Sandia National Laboratories. We then applied all this technology to perform simulations on the heterogeneous POWER9 cluster installed at the Barcelona Supercomputing Center, with an architecture very similar to that of the Summit supercomputer from the Oak Ridge National Laboratory – ranked first in the top500 list. In the BSC POWER9 cluster, which has 4 NVIDIA P100 GPUS per node, we demonstrated that we could perform a well-balanced co-execution using both the CPUs and GPUs simultaneously, being 23% faster than using only the GPUs. In practice, this represents a performance boost equivalent to attaching an additional GPU per node. This research was published at the Future Generation Computer Systems journal [2] where a full analysis of the code performance is given. Sample results of the elapsed time per MPI Rank is given in Figure 1 (left), while a snaphsot of Q-vorticity along the wing is shown in Figure 1 (right).

References: [1] R. Borrell, G Oyarzun, et al. Proceedings of ScalA 2019: 10th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems - Held in conjunction with SC 2019: The International Conference for High Performance Computing, Networking, Storage and Analysis, 8948628, pp. 72-78, 2019. [2] R. Borrell, D. Dosimont et al. Heterogeneous CPU/GPU co-execution of CFD simulations on the POWER9 architecture: Application to airplane aerodynamics. Future Generation Computer Systems, (107): 31-4, 2020.

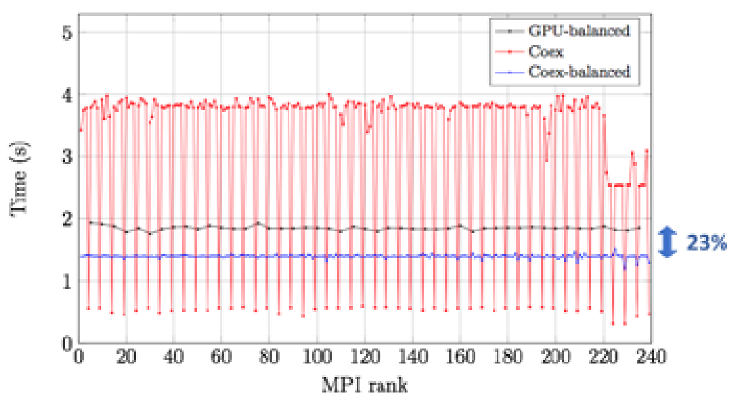

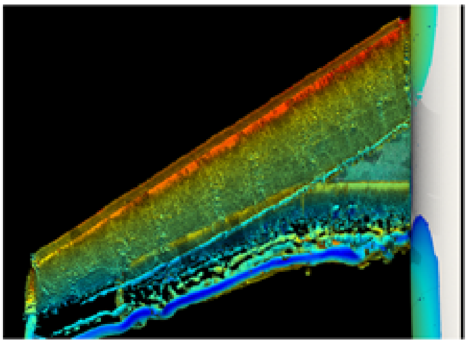

Figure 1. (Left): Comparison of (balanced) co-execution vs pure GPU execution – elapsed time per MPI Rank. (Right): Snapshot of Q iso-surfaces of the turbulent flow around an airplane.

Enabling Exascale in the CFD code AVBP 1/2

Computational fluid dynamics remains an ever complex challenge. The increase in computational resources has allowed incredible strides in numerical modelling towards fully digital understanding of complex flows. Within EXCELLERAT, AVBP is one such application dedicated to combustion. AVBP is a thirty-year-old computational fluid dynamics code used by research and industry alike to compute the 3D Navier-Stokes reactive two phase equations in a variety of applications from experimental test-beds, piston engines and up to rocket engines. These technical challenges are computationally very complex and require massive resources. This is where high-performance computing and exascale intervene.

Two use cases have been singled out to showcase the capabilities and potential of exascale in combustion: first, a safety application of a modelled explosion inside a confined space and second, a pollutant emissions case in an aeronautical type configuration.

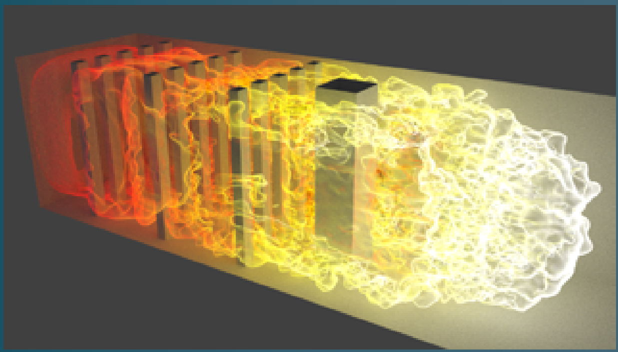

Fig. 1: Instantaneous view of a temperature iso-surface (textured with blender). Gullaud et al CERFACS

Both these cases incorporate complex multi-scale problems (from millimeters for the flame thickness to meters for the wall distances for example) that can only run on Tier-0 systems like the CINECA or TGCC systems from PRACE.

Within EXCELLERAT, we are tackling the challenges towards the next generation Tier-0 and above systems to perform these simulations even more efficiently. First, by introducing an automatic mesh refinement framework within the AVBP code and also by accelerating this legacy code using Graphical Processing Units (GPU) acceleration. To tackle AMR, a refactoring of the data structure was undertaken to allow for dynamic meshes at runtime. A packed data structure for the grid was introduced by the CFD Team at CERFACS. The runtime AMR approach consists of a phase where the computation is stopped, adaptation and load balancing are offloaded to a library and then the work resumes after reconstructing the parallel communication information.

In parallel, two approaches have been tested for mesh refinement. The first one is a collaboration between CERFACS and CORIA-CNRS using their YALES2 code to perform the mesh adaptation, coupling it with AVBP. Also an in-house tool TREEADAPT, developed on top of the TREEPART load balancing library is being tested (figure 2 shows the different phases of

Enabling Exascale in the CFD code AVBP 2/2

adaptation using TREEADAPT, a first adaptation is performed freezing the parallel interfaces, then a second load balancing and adaptation is performed to deal with the remaining untouched elements).

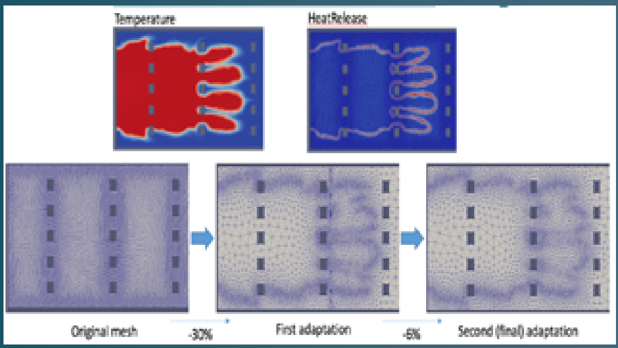

Fig. 2: mesh adaptation phases using TREEADAPT

So far, the AVBP-YALES2 method has been tested successfully on a 2D explosion case yielding a 2x acceleration compared to the fully resolved case, a 3D case starting with a coarse mesh and tests with the TREEADAPT method are underway.

Furthermore, recent advances on parallel programming runtimes have rendered GPU programming within reach of traditional computing codes. For AVBP, a Fortran code, the OpenACC framework was chosen for its simplicity compared to OpenMP (regarding memory management). The fruit of a collaboration with GENCI, IDRIS, initially IBM and then HPE, OpenACC has been gradually introduced at the loop level following the fine-grain approach. This progressive and methodic albeit slow approach allows for systematic code verification and validation of each kernel adding more than 3,000 pragmas.

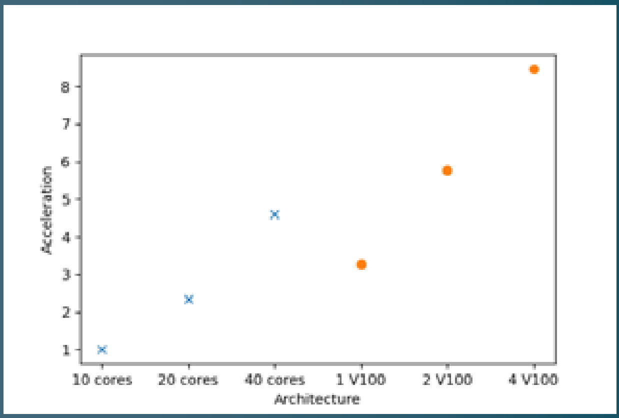

Fig. 3: Acceleration for a 3D explosion simulation without AMR using full MPI (X) and MPI+OpenACC (O) and a 40 cores Intel Cascade lake system equipped with 4 V100 Nvidia GPU (Jean Zay system at IDRIS from GENCI).

This first port of AVBP to GPU already shows an 8x acceleration on 4 V100s compared to 10 cascade lake cores. Strong scaling has been tested on the PIZ DAINT CSCS/PRACE system equipped with P100 GPUs. Scaling has been tested up to 32 GPUs with encouraging results (figure 4), negative impact on performance of cache blocking techniques developed for CPUs, as well as the need for further optimisation.

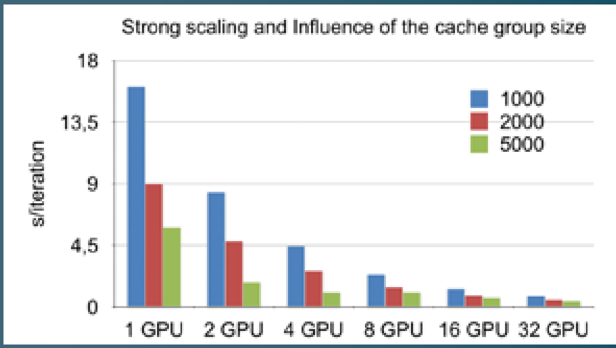

Fig. 4: Strong scaling of AVBP on Piz Daint P100s and influence of the cache blocking on performance.

Nek5000 in EXCELLERAT: a use case

Nek5000 is one of the reference codes selected for EXCELLERAT. It is an academic high-order CFD solver based on a Spectral Element Method used for modelling a number of different flow cases including industrial application e.g. nuclear reactor thermal-hydraulics. An important improvement to the code was the implementation of an Adaptive Mesh Refinement (AMR) algorithm that allows control of computational error during the simulation and to reduce it at minimal cost by proper mesh adjustment. This code is currently extended within EXCELLERAT for fully three-dimensional flows making it a robust solver capable of solving industrially relevant problems. A number of different aspects of numerical modelling are considered here starting from mesh generation, solver parallel performance, through solution reliability (Uncertainty Quantification) and ending with in-situ data analysis and visualisation. Developed tools are tested with simulation of a flow around three-dimensional wing tip, which was selected as one of use cases for EXCELLERAT. It is a relatively complex test case relevant for a number of industrial areas like aeronautics, automotive or bio-engineering. However, despite its environmental and economic importance, there is a lack of high-fidelity numerical data for high-Reynolds-number turbulent flows.

We started with creating a workflow for the pre-processing step including hex-based mesh generation with widely used meshing software gmsh, and development of proper tools for the simulation initialisation. It was used to generate a suitable for Nek5000 mesh of NACA0012 aerofoil with a rounded wing tip. An important aspect here was the proper definition of relative element orientations, which is necessary to define global gather-scatter operations in the solver. We focused as well on a projection of the grid points located at the domain boundary on the prescribed wing surface. This operation is required by the AMR algorithm, as mesh modifications using simple interpolation lead to corrugation of the boundaries. However, it is a challenging operation for the wing tip test case, as it contains a sharp trailing edge introducing in our parametrisation a singular point. Mesh refinement in this case can give the invalid elements that have to be corrected by proper shift of the grid points at the boundary using spectral interpolation before projection.

All the described modifications enabled us to perform full AMR simulations of the proposed use case. Stay tuned for more information about Nek5000 towards exascale in our next issue.

N. Jansson, A. Peplinski, S. Rezaeiravesh, M. Zavala

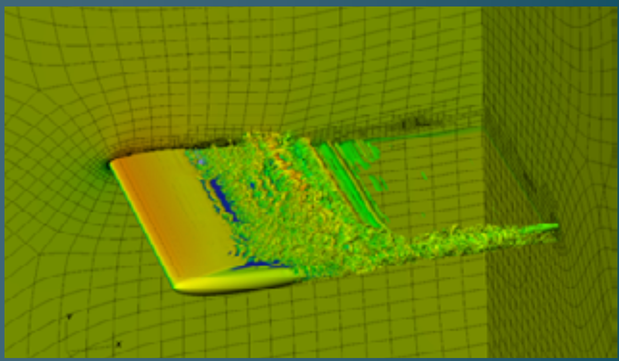

The figure presents a vortical structure of the transient flow around NACA0012 aerofoil with rounded tip for Re=105. Colour scale gives a chord wise velocity component and black lines mark the element boundaries showing the regions with increased resolution.

News

Podcast

We launched a by-monthly HPC podcast powered by EXCELLERAT, with expert talks about exascale news, discussions and engineering issues to be solved by exascale computers. Its first episode deals with the topic „What is HPC all about?“. In this expert interview, Andreas Wierse, Managing Director of SICOS GmbH, explains what High Performance Computing is and what it can be used for in the engineering area.

Stay tuned for the next episodes:
https://soundcloud.com/user-61965686

HPC Trivia Quiz

On Twitter, we launched a trivia quiz series powered by EXCELLERAT with a focus on HPC topics. New questions will be released every Monday and you can find the answers on each Wednesday afterwards.

Join the quiz:
https://twitter.com/EXCELLERAT_CoE

All Hands Meeting (5-7 May)

Due to the travel restrictions linked to corona virus, we have organised our first EXCELLERAT virtual All Hands Meeting. Instead of hosting it in Stockholm from 5th until 7th May, so far this was our team's largest online conference. During this All Hands Meeting, we presented the status quo of our project's work packages, discussed about our codes' use cases, exchanged knowledge and experience on behalf of our working groups, and prepared our upcoming project review.

www.excellerat.eu

@EXCELLERAT_CoE EXCELLERAT

The EXCELLERAT project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823691