Joint EERA–EoCoE Position Paper on HPC for Energy

In response to the EU's policy agenda for the energy sector, which links the strategy to accelerate decarbonization with the adoption of digital technologies, the European Energy Research Alliance (EERA) together with the Energy-oriented Centre of Excellence (EoCoE) has presented a joint position paper to the European Commission on the subject of exascale, and the great opportunity it represents for the clean energy transition in Europe.

As Europe works towards a decarbonized energy ecosystem, with a clear vision and goals set by the European Green Deal, EERA and EoCoE have identified a blind spot: energy domain scientists do not take full advantage of the potential that HPC-fuelled simulations can offer for their work. This situation is the result of a lack of HPC-related expertise available to scientists.

To this end, EERA recently created a transversal Joint Programme entitled "Digitalisation for Energy", which allows a wide range of scientists in the energy domain to access the considerable expertise amassed through previous cross-domain collaborations between application experts and domain scientists explored in the EoCoE project. A number of pilot collaborations have already been established via this Joint Programme (in materials modelling, hydropower, and energy systems integration), with more to follow.

The position paper aims to trigger suitable coordination actions and funding decisions from the European Commission and Member States to support the development of tuned data models and simulation codes for topics in the field of energy. To do so, the paper proposes to make use of the latest technology in high-performance computing (HPC) and data management that is becoming available at EU level: the exascale generation. The paper was reviewed by a scientific committee including representatives from Forschungszentrum Jülich, and the editorial tasks were coordinated by JSC. The paper explores the role played by digital tools tuned for the energy sector in support of the transition towards climate neutrality. The full paper can be downloaded from the EERA website.

Contact: Dr. Maria Ramalho, m.ramalho@fz-juelich.de

JSC@ISC2022

The conference ISC High Performance 2022 will take place from 29 May to 2 June 2022 in Hamburg, Germany (https://www.isc-hpc.com).

JSC, together with its partners in the Gauss Centre for Supercomputing (GCS) – HLRS in Stuttgart and LRZ in Garching – will present its wide-ranging supercomputing activities at the GCS booth. Focus topics include the path to exascale with JSC's Modular Supercomputing Architecture (MSA) concept; the steadily increasing importance of AI in HPC, reflected by developments such as the Helmholtz AI Cooperation Unit; and quantum computing technologies, where the growing Jülich UNified Infrastructure for Quantum computing (JUNIQ) will be presented.

JSC employees will also contribute to the event with numerous talks, tutorials, and workshops. For example, Andreas Herten is co-organizing the tutorial "Efficient Distributed GPU Programming for Exascale"; Markus Geimer will be a speaker at the tutorial "Practical Hybrid Parallel Application Performance Engineering"; Sebastian Achilles will present at the tutorial "Maintaining a Modern Scientific Software Stack Made Easy with EasyBuild"; Bernd Mohr and Morris Riedel are organizing the tutorial "Introduction to HPC Applications Systems, Programming Models and Machine Learning and Data Analytics"; Bernd Mohr will chair the Research Paper Session "Productivity Tools and Performance Modeling & Tuning"; Stefan Kesselheim will discuss "MLPerf: A Benchmark for Machine Learning" at a Bird-of-a-Feather session; and Estela Suarez is organizer of the workshop "Malleability Techniques Applications in High-Performance Computing".

Detailed information on JSC's participation and activities can be found at https://go.fzj.de/isc22.

Contact: Jens Henrik Göbbert, j.goebbert@fz-juelich.de

JURECA Booster to Be Decommissioned in October

JSC will decommission the JURECA Booster at the end of October 2022. In November 2021, 400 nodes were already taken offline to provide an additional power budget for JURECA-DC. The remaining 1,240 nodes will stay available until the end of October. Because the JURECA Booster is accessed via JURECA-DC logins, the access pattern will not change. Users who currently rely on the JURECA Booster should migrate their workloads to other systems like JURECA-DC or the JUWELS Booster, given the additional performance these platforms provide.

The JURECA Booster entered production in 2017 and extended the existing JURECA Cluster supercomputer to create the first production-ready modular supercomputer, a result of the DEEP projects and prototypes that paved the way for the development of the Modular Supercomputing Architecture (MSA). At the time, JURECA was the first system to combine high-speed NVIDIA Mellanox InfiniBand and Intel OmniPath interconnects via bridging on the hardware and software level. The deployment involved a concerted effort by the participating partners, namely Dell, Intel, JSC, ParTec, and T-Platforms. JURECA was the first modular supercomputer to be included on the Top500 list. In November 2017, it was ranked 30th with the Linpack benchmark running across both the JURECA Cluster, installed in 2015 with 2.3 petaflops, and the JURECA Booster, with a peak performance of 5 petaflops based on Intel Knights Landing processors.

Contact: Benedikt von St. Vieth,
b.von.st.vieth@fz-juelich.de

HPST Enters Full Production

The High Performance Storage Tier (HPST) has entered full production and is opening access to more projects via JUDOOR. It is available as $CSCRATCH. Acting as a NVMe-based cache layer on top of JUST, it offers faster access to data stored in the SCRATCH file system. This added storage sub-system further closes the gap between the GPFS file system on one hand and the compute systems on the other. Data analysis and I/O intensive applications, in particular read intensive AI workloads, can greatly benefit from staging data to a faster storage tier.

To achieve high bandwidth and low latency the HPST is directly integrated into the Infiniband Fabric of all 3 main systems JUWELS, JURECA and JUSUF, allowing for a peak bandwidth of up to 2TB/s. To minimize efforts from applications to utilise the system, the underlying 2 PB of SSD storage space can be directly accessed using POSIX. However, for maximum performance it is recommended to use the ParaStationMPI MPI-IO or the native IME interface.

Prior to launching their respective jobs on a compute system, users can use the available IME tools to stage data from the back-end SCRATCH file system to the HPST and after job end synchronise it back. Using this strategy, some applications have observed an I/O time reduction of up to 5x during job execution. The support team checks the I/O pattern of applications requesting access to the HPST to maximise overall system benefit. Respective candidates are assigned a quota on the HPST, which when exceeded data will be automatically evicted, allowing for parallel use by multiple applications without interference between users. The HPST system has been supplied by HPE and DDN, and runs the DDN developed Infinite Memory Engine (IME). The procurement is part of the Fenix e-infrastructure within the scope of the EU project ICEI.

Contact: Dr. Salem El Sayed Mohammed, s.elsayed@fz-juelich.de

Google Summer of Code

For the first time, Forschungszentrum Jülich will be participating in Google Summer of Code as a mentoring organization with the JSC-led project Heat. Initially funded within the Helmholtz Analytics Framework (a Helmholtz Information & Data Science pilot project), Heat is an open-source Python library developed with the goal of enabling data-intensive research – multi-CPU and multi-GPU computing, and highly optimized machine learning and deep learning – using a NumPy/SciPy-compatible API. Its general, easy-to-adopt solutions to single-node memory bottlenecks have led to an expansion of interest and feature requests from outside the Helmholtz Association.

Google Summer of Code (GSoC) is a worldwide online programme that aims to introduce new contributors to open-source software development. Over the course of the summer, GSoC contributors mentored by members of the Heat core team (from JSC, DLR, and KIT) will work on implementing new features to support our new and established user communities, from memory-distributed signal processing to distributed sparse operations, parallel dimensionality reduction algorithms, expanded support of hardware accelerators, and more.

Interested participants are already contributing to Heat's repository, engaging on our Mattermost channel, and drafting proposals. Coding officially starts on 14 June, preceded by a three-week Community Bonding Period. We relish the opportunity to work with open-source enthusiasts, support our expanding user base, and act as pathfinders for future GSoC participation by other open-source projects under the FZJ umbrella. More information about Heat can be found on GitHub:
https://github.com/helmholtz-analytics/heat

Contact: Dr. Claudia Comito, c.comito@fz-juelich.de

Events

Introduction to Supercomputing at JSC – Theory & Practice

Instructors: JSC employees, representatives of Atos, Intel, and ParTec
Date: 16–19 May 2022, starting on 16 May at 13:00
Venue: online
https://go.fzj.de/2022-sc-1