

9th bwHPC Symposium

Monday, October 23, 2023 - Monday, October 23, 2023

University of Mannheim



Book of Abstracts

Contents

Welcome from the rectorate of the University of Mannheim	1
Welcome from MWK	1
Machine Learning and Simulation - Challenge and Chance for HPC	1
NHR: HPC resources for your science	1
Base4NFDI - Creating NFDI-wide basic services in a world of specific domains	1
bwHPC-S5	2
A Fully Parallelized and Budgeted Multi-Level Monte Carlo Method and the Application to Acoustic Waves	2
Simulation of pyrolysis process of waste plastics	2
Orchestrating Gaia-X Workflows in Cluster-Environments	3
Analysis of one photon calcium imaging data to investigate the properties of CA1 place cells	3
cellular_raza - Novel Flexibility in Design of Agent-Based Models in Cellular Systems	4
ABI Checking for Libraries and Applications	4
How To GDPR? Handling Sensitive Data on bwForCluster BinAC 2	5
3D general relativistic collapse of rotating neutron stars into black holes	5
Concurrent multi-scale simulation of lubrication, enabled by machine learning	6
Exploring the limits of new physics at the LHC	6
End-user data analysis at the LHC	6
Fair sharing of resources between clusters with AUDITOR	7
Propane/Propylene Separation via Molecular Dynamics Simulations in Metal-Organic Frameworks.	7
Computational Microstructure Exploration for Batteries: Leveraging High Performance Computing	8
Advanced Training-Data Generation for Machine-Learned Potentials	8

bwFDM –The Federal State Initiative for Research Data Management in Baden-Württemberg	9
Base4NFDI - Basic Services for NFDI	9
Further development of the state service "Scientific Data Storage"	10
bwForCluster Helix	10
Introducing Galaxy: A new cloud and HPC gateway for data-driven computational research	10
The agroecosystem multi-model library, Expert-N on high-performance computers (HPC)	11
High-resolution global weather simulations with the Model for Prediction Across Scales (MPAS) on the Hawk HPC system	11

Welcome Speeches / 156**Welcome from the rectorate of the University of Mannheim**

We welcome you to the University of Mannheim for the 9th bwHPC-Symposium!

Welcome Speeches / 160**Welcome from MWK****Keynote Speech / 172****Machine Learning and Simulation - Challenge and Chance for HPC**

Author: Heiner Stuckenschmidt¹

¹ *University of Mannheim*

Corresponding Author: heiner.stuckenschmidt@uni-mannheim.de

In recent years, users have been asking more and more for support with large-scale machine learning applications leading to a focus on memory and GPU computation. In this talk I argue that the next big thing will be the combination of machine learning and simulation and discuss a number of applications based on this combination and the computational challenges that result from this combination.

bwHPC, NHR and NFDI / 183**NHR: HPC resources for your science****bwHPC, NHR and NFDI / 180****Base4NFDI - Creating NFDI-wide basic services in a world of specific domains**

Author: Franziska Fritzsche¹

¹ *Gesis - Leibniz-Institute for Social Sciences*

Corresponding Author: franziska.fritzsche@gesis.org

Today, research data is often stored in many different places, difficult to find and only available for a limited time. Base4NFDI as a joint initiative of all consortia of the German National Research Data Infrastructure (NFDI) aims to create the basis for better findability, accessibility, interoperability and reusability of research data. For this purpose, common, technical services are developed together with experts for the data in the different research disciplines. Since many scientific fields have

similar requirements for research data, Base4NFDI supports common solutions to avoid parallel developments. Already existing services are thereby adapted or extended to be usable for researchers from other disciplines.

All institutions in the scientific community (and beyond) can actively participate in the development and implementation of these services and submit proposals via the sections of NFDI. A panel of experts selects the proposals of the development teams for Base4NFDI, which then receive funding from the German Research Foundation (DFG). These teams also receive support and guidance in areas such as development, implementation, and training. After development, the NFDI basic services will be offered permanently to the scientific community. Base4NFDI thus actively contributes to the systematic opening and networking of the German science system.

The presentation will give an insight into the idea and structure of Base4NFDI, the processes for developing basic services and the services currently funded.

bwHPC, NHR and NFDI / 127

bwHPC-S5

Corresponding Author: robert.barthel@kit.edu

The presentation outlines the current status of the user support project bwHPC-S5 for HPC, DIC and LS2DM in the state of Baden-Württemberg, Germany.

The state-wide project bwHPC-S5 provides federated support for the bwHPC users and coordinates all associated provisions, including development of policies and services. As the connecting body between scientists and HPC systems, bwHPC-S5 has proven to create synergies for the development of state-wide user support. The implementation of professional competence centers guarantees the support expertise required to embed the scientific communities into the HPC, DIC and LS2DM world and to increase efficiency and effectiveness in utilizing both, the compute and storage resources, by optimizing workflows as well as performance and scalability of applications. bwHPC-S5 extends the federated services established in previous projects to include the acquisition, processing, storage and archiving of large scientific datasets.

Talks / 163

A Fully Parallelized and Budgeted Multi-Level Monte Carlo Method and the Application to Acoustic Waves

Authors: Niklas Baumgarten^{None}; Christian Wieners^{None}; Sebastian Krumscheid^{None}

Corresponding Authors: sebastian.krumscheid@kit.edu, niklas.baumgarten@kit.edu

We present a novel variant of the multi-level Monte Carlo method that effectively utilizes a reserved computational budget on a high-performance computing system to minimize the mean squared error. Our approach combines concepts of the continuation multi-level Monte Carlo method with dynamic programming techniques following Bellman's optimality principle, and a new parallelization strategy based on a single distributed data structure. Additionally, we establish a theoretical bound on the error reduction on a parallel computing cluster and provide empirical evidence that the proposed method adheres to this bound. We implement, test, and benchmark the approach on computationally demanding problems, focusing on its application to acoustic wave propagation in high-dimensional random media.

Talks / 171

Simulation of pyrolysis process of waste plastics

Author: Feichi Zhang¹

Co-authors: Salar Tavakkol¹; Thorsten Zirwes²; Dieter Stapf¹

¹ *KIT/ITC*

² *ITV/Uni Stuttgart*

Recycling of waste represents one of the major goals to close the carbon cycle and to achieve circular economy concepts. The thermal conversion of plastic waste via pyrolysis to base chemicals for the production of new plastics represents one of the key technologies.

The talk provides an overview of fundamental research carried out at the Institute for Technical Chemistry (ITC) at KIT, with the objective to develop and optimize high-temperature process techniques for plastics pyrolysis suited for industrial-scale applications. Particular focus of the work is to study the plastics pyrolysis process by means of high-fidelity numerical simulations, which can be regarded as a “digital twins” of the real systems. In this way, the simulations provide detailed information for the temporal evolution and spatial distributions of the flow and chemical scalar fields, such as pressure, velocity, and temperature or species concentrations, which often cannot be assessed experimentally. Moreover, the modeling concept enables detailed analyses of the mutual interactions between underlying chemical-physical processes, such as heat transfer and chemical reactions, which dominate the pyrolysis process of plastics.

A couple of ongoing studies will be presented, ranging systematically from basic setups using a single plastic particle to a fully captured laboratory-scale fluidized bed reactor, which covers millions of sand and plastic particles. The results reveal strong correlations between characteristic behaviors of the pyrolysis process with the applied operating parameters. These correlations can be used to develop a more efficient reactor design for the pyrolysis of plastic waste. The opportunities and challenges for applications of the modeling concept concerning a compromise between simulation accuracy and computational cost, as well as the predictability of the underlying chemo-physical effects, e.g. pyrolysis reactions and melting behavior, will be highlighted at the end.

Poster Session / 170

Orchestrating Gaia-X Workflows in Cluster-Environments

Authors: Ralph Isenmann¹; Janis Keuper¹

¹ *HS Offenburg*

Corresponding Author: janis.keuper@hs-offenburg.de

Gaia-X is an initiative that develops, based on European values, a digital governance that can be applied to any existing cloud/ edge technology stack to obtain transparency, controllability, portability and interoperability across data and services. In this poster, we will present the intermediate results from our ongoing BMWk funded **Gaia-X4AI** project (<https://gaia-x4ki.eu/>), developing scalable gaia-x workflows for AI applications on compute clusters. We introduce our novel middle-layer, which allows to abstract resource aware scheduling of containerized gaia-x workflows on top of Slurm based HPC systems as well as on Kubernetes Clusters. We evaluate our approach on Use Cases in the context of autonomous driving applications, e.g. showing the deployment of compute intensive CARLA based simulations on our technology stack.

Poster Session / 177

Analysis of one photon calcium imaging data to investigate the properties of CA1 place cells

Author: Vladislav Ivantaev^{None}

Co-authors: Alessio Attardo¹; Christian Leibold²

¹ *Leibniz Institute for Neurobiology*

² *Albert-Ludwigs-Universität Freiburg, Fakultät für Biologie & Bernstein Center Freiburg*

Corresponding Author: vladislav.ivantaev@bcf.uni-freiburg.de

The hippocampus is considered to be involved in spatial navigation and memory formation; it hosts cells with high firing rates at a given location –place cells (PCs), which allow the construction of the hippocampal cognitive map of space. However, it is not understood how memories are related to the hippocampal place fields. On the one hand, the stability of its neural code may support long-term storage and retrieval of memories; on the other hand, the turnover of its neural representations may facilitate learning and adaptation.

Here, we focus on the analysis of one-photon (1P) calcium imaging data from dorsal CA1 of a mouse, recorded with a miniature microscope during free foraging in a previously familiarized two-dimensional arena, while most of the calcium imaging data from the hippocampus obtained from one-dimensional tracks. The CaImAn software tool is applied to the raw recordings to extract the spatial footprints of cells, as well as the corresponding time-varying calcium transients, which are used to estimate the cells' place fields and spatial information (SI) values. With the CellReg toolbox, neurons were tracked across multiple sessions, allowing to investigation of the temporal dynamics of CA1 neural code over time on a scale of up to two weeks.

After doing that three methods of determining statistically significant PCs are benchmarked on the aforementioned set of data comprising 5-6 sessions for 2 subjects.

Poster Session / 181

cellular_raza - Novel Flexibility in Design of Agent-Based Models in Cellular Systems

Authors: Jonas Pleyer^{None}; Christian Fleck^{None}

Corresponding Author: jonas.sci@pleyer.org

Agent-Based Models (ABMs) allow researchers to describe complex cellular systems in a mechanistic manner but can also abstract over less-known processes.

It is often desirable to exchange only parts of the model eg. changing the spatial representation of cells from a spherical interaction potential to an elliptical.

Existing tools lack in flexibility and cannot change their internal representation of cells.

To solve these problems we created cellular_raza, a novel library that offers previously unknown flexibility in model design while retaining excellent performance.

Poster Session / 185

ABI Checking for Libraries and Applications

Author: Rainer Keller¹

¹ *HS Esslingen*

Corresponding Author: rainer.keller@hs-esslingen.de

Application Binary Interface (ABI) compliance is important for libraries to adhere to. Open MPI, being used by many applications and using several vendor libraries itself, is an important library by itself.

This poster provides methods to check for ABI compliance in new installations and best practices for library developers to make sure that they adhere to their user's expected interface – or work with semantic versioning to make sure that applications are updated.

Poster Session / 186

How To GDPR? Handling Sensitive Data on bwForCluster BinAC 2

Authors: Stefan Behnle¹; Rafael Doros²; Volker Lutz²; Jens Krüger²; Felix Bartusch²

¹ *Eberhard Karls Universität Tübingen*

² *Eberhard Karls Universität Tübingen*

Corresponding Authors: felix.bartusch@uni-tuebingen.de, rafael.doros@uni-tuebingen.de, stefan.behnle@uni-tuebingen.de, jens.krueger@uni-tuebingen.de, volker.lutz@uni-tuebingen.de

Modern experimental techniques and medical examination methods generate vast amounts of data that must be protected carefully according to the GDPR. At the same time, handling of such data, e.g., human DNA sequences, MRI images, or EEG traces is necessary for medical research and personalized medicine. Especially when performing statistical analyses on many datasets, the amount of data is too big to be handled on personal computers or local workstations.

High Performance Computing (HPC) clusters are ideally suited for handling these tasks even on the terabyte scale in acceptable time. Their typical architecture however is optimized towards performance and less towards data security.

We present concepts on how to provide users of the bwForCluster BinAC 2 with the computational power of a supercomputer as well data integrity and security compliant with the GDPR and other data protection regulations. This can be achieved by separation and isolation of storage, network, and compute nodes for the handling of sensitive data. Dedicated techniques cover the creation of isolated on-the-fly subclusters, use of advanced containerization and virtualization technologies and extensive logging for detecting and recording unauthorized data access. At the heart of these activities is the definition and documentation of procedures jointly with the users of the system. The overarching goal is the establishment of an information security management system and its certification according to ISO27001.

Poster Session / 187

3D general relativistic collapse of rotating neutron stars into black holes

Author: José Carlos Olvera Meneses¹

Co-author: Daniela Doneva¹

¹ *Tübingen University*

Corresponding Author: jose.olvera-meneses@uni-tuebingen.de

In this work, we investigate the phenomenon of neutron star collapse into a black hole within the framework of modified theories of gravity, exploring the consequences of departures from General Relativity (GR), specifically, under massive scalar-tensor theories that allow for scalarization in order to understand the effect of these modifications in the astrophysical process. To accomplish this, we employ advanced numerical techniques and high-performance computing to accurately model the dynamics of the star's core as it approaches the critical density for collapse and the posterior

stage when the black hole is formed. We will also compute the gravitational radiation and compare our findings with observational data to constrain the parameters of these alternative gravity theories.

Talks / 168

Concurrent multi-scale simulation of lubrication, enabled by machine learning

Authors: Hannes Holey¹; Mohamed Tarek Elewa¹; Gumbsch Peter²; Lars Pastewka³

¹ *Karlsruhe Institute of Technology*

² *Fraunhofer IWM*

³ *University of Freiburg*

Corresponding Author: lars.pastewka@imtek.uni-freiburg.de

Friction and lubrication are inherent multiscale problems, particularly when the gap between contacting bodies is on the order of molecular interaction length scales, such as in the boundary lubrication regime. Modelling lubrication across scales beyond purely sequential approaches has so far remained elusive. In this talk, I will present a reformulation of the classical lubrication equations that principally allows straightforward coupling between continuum and molecular models. Concurrency is achieved by informing a surrogate model for the constitutive behavior of highly confined fluids on-the-fly using molecular dynamics (MD) simulations. An active learning scheme based on Gaussian process regression allows data-efficient interpolation of microscopic stresses obtained from MD in possibly high-dimensional parameter spaces. The proposed method is validated for simple fluids before application to more realistic lubricant models.

Talks / 169

Exploring the limits of new physics at the LHC

Author: Nikita Schmal¹

¹ *Heidelberg University*

Corresponding Author: schmal@thphys.uni-heidelberg.de

The vast amount of data produced at the Large Hadron Collider (LHC), the largest particle collider in the world, is used to study the most fundamental building blocks of nature. Global analyses using LHC data provide ways to constrain signs of new physics within these measurements. In this talk, we present an overview of the uses of the NEMO cluster in these analyses, utilizing its strength to simulate collider events and explore high-dimensional parameter spaces.

Talks / 176

End-user data analysis at the LHC

Authors: Nils Faltermann^{None}; Günter Quast¹; Manuel Giffels¹

¹ *KIT*

Corresponding Authors: nils.faltermann@kit.edu, guenter.quast@kit.edu, manuel.giffels@kit.edu

The Large Hadron Collider (LHC), located at CERN in Geneva, stands as one of the most monumental scientific experiments in human history. This remarkable machine facilitates approximately 40 million particle collisions every second, generating an astronomical amount of data. Even with rigorous filtering of collision events, the data retained for subsequent analysis remains staggering in scale.

In addition to the recorded data, conducting a successful physics analysis demands an extensive set of simulations that can be compared to the recorded events. In this presentation, we will delve into our approach to incorporating external resources, such as the NEMO cluster in Freiburg, into our local batch system. This integration greatly enhances accessibility for the complex workflows required for physics data analyses.

Talks / 173

Fair sharing of resources between clusters with AUDITOR

Authors: Benjamin Rottler¹; Michael Boehler¹; Anton J. Gamel¹; Dirk Sammel¹; Markus Schumacher¹

¹ *Albert-Ludwigs-Universität Freiburg*

Corresponding Authors: michael.boehler@physik.uni-freiburg.de, anton.gamel@physik.uni-freiburg.de, dirk.sammel@physik.uni-freiburg.de, markus.schumacher@physik.uni-freiburg.de, benjamin.rottler@physik.uni-freiburg.de

For several years, we have been dynamically and opportunistically integrating the computing resources of the HPC cluster NEMO into the HTC cluster ATLAS-BFG using the COBalD/TARDIS software. To increase usage efficiency, we allow the integrated resources to be shared between the various High Energy Physics (HEP) research groups in Freiburg. However, resource sharing also requires accounting. This is done with AUDITOR (AccoUnting DatahandlIng Toolbox for Opportunistic Resources), a flexible and extensible accounting ecosystem that can cover a wide range of use cases and infrastructures. Accounting data is recorded via so-called collectors and stored in a database. So-called plugins can access the data and take measures based on the accounting documents. In this work, we present how NEMO resources can be fairly shared among contributing working groups when integrated into ATLAS-BFG using AUDITOR.

Talks / 164

Propane/Propylene Separation via Molecular Dynamics Simulations in Metal-Organic Frameworks.

Author: Mikhail Suyetin^{None}

Corresponding Author: msuyetin@gmail.com

Propylene is crucial for the petrochemical industry, and most studies focus on separating it from propane. However, current methods for obtaining pure propylene require complex, energy-intensive desorption processes using propylene-selective porous materials. A more efficient approach is to develop an adsorbent that prefers propane, allowing for one-step high-purity propylene production, reducing energy consumption and the need for large amounts of adsorbent. To find suitable materials, the CoRE MOF database is used, employing molecular dynamics simulations in identifying materials with strong potential for propane/propylene separation. The top 5 MOFs for propane and propylene separation are identified, with carbonyl groups found to significantly enhance separation. Additionally, a machine learning model is used to predict self-diffusion values, showing good agreement with molecular dynamics simulation data.

Talks / 166

Computational Microstructure Exploration for Batteries: Leveraging High Performance Computing

Author: Julius Weinmiller¹**Co-authors:** Benjamin Kellers¹; Martin P. Lautenschläger¹; Timo Danner¹; Arnulf Latz¹¹ *Deutsches Zentrum für Luft- und Raumfahrt (DLR) e. V.***Corresponding Authors:** arnulf.latz@dlr.de, benjamin.kellers@dlr.de, timo.danner@dlr.de, julius.weinmiller@dlr.de, martin.lautenschlaeger@dlr.de

Battery performance is strongly influenced by physical and electrochemical processes occurring on the pore scale. Thus, especially the microstructure of battery components is important. Therefore, battery research is highly interested in understanding the interdependence of both. Our work goes into this direction applying computational methods that are used in the field of battery research only recently. Those are the lattice Boltzmann method (LBM) and pore network modeling (PNM). Both are further developed and used in a smart and complementary manner. The application of which is presented for two recent research topics: 1) the influence of chemical surface reactions on transient battery morphologies, and 2) the identification of representative elementary volumes for reducing computational cost. Both topics are placed in the context of our efforts, i.e., applying high performance computing tools for developing and understanding battery technologies.

Talks / 167

Advanced Training-Data Generation for Machine-Learned Potentials

Authors: Moritz René Schäfer¹; Johannes Kästner¹¹ *University of Stuttgart***Corresponding Authors:** kaestner@theochem.uni-stuttgart.de, schaefer@theochem.uni-stuttgart.de

Molecular dynamics (MD) plays a crucial role in the field of atomistic simulations for calculating thermodynamic and kinetic quantities of molecules and materials.

Traditionally descriptions of atomic interactions either rely on classical or quantum mechanical models which are limited in accuracy and simulation speed respectively.

Machine learning potentials (MLPs) have emerged as a useful class of surrogate models that bridge this gap, retaining most of the quantum mechanical accuracy at drastically reduced cost.

However, obtaining informative training data for these models is a challenging task, as the systems of interest can have thousands of degrees of freedom with vastly different characteristic time scales.

Biasing the dynamics of the system along the slowest degrees of freedom can significantly decrease the time needed to obtain sufficiently informative data for the training of MLPs.

While traditional approaches to biasing MD mostly rely on hand-selected degrees of freedom to enhance the sampling in, we introduce a data-driven way to identify the most informative degrees of freedom for the MLP and limit the bias to exploring physically relevant parts of configuration space.

Further, efficient execution of atomistic machine learning workflows relies on the utilization of heterogeneous compute resources.

While model training and MD simulations are most efficient on GPUs, reference quantum mechanical computations require large amounts of CPU cores.

We introduce tools to flexibly split the tasks in a workflow across the available hardware.

Poster Session / 178

bwFDM – The Federal State Initiative for Research Data Management in Baden-Württemberg

Authors: Cora F. Krömer¹; Fabian Schubö¹

¹ Karlsruhe Institute of Technology

Corresponding Authors: cora.kroemer@kit.edu, fabian.schuboe@kit.edu

bwFDM is engaged in various areas of research data management (RDM). The main focus is on establishing networks between stakeholders, providing information on RDM-related topics, offering consultations and training, and hosting the biennial conference series *E-Science-Tage*. Several of the services offered by *bwFDM* are relevant to research groups, *bwHPC* service providers, and *bwHPC* support centers. A central aim is to connect these stakeholders with RDM consultants and other stakeholders from the research data infrastructure in Baden-Württemberg. Therefore, *bwFDM* is in constant exchange with *bwHPC-S5* on cross-cutting topics.

A major *bwFDM*-flagship is *forschungsdaten.info*, the leading information platform for RDM in Germany. The platform provides an introduction and overview of RDM topics, services and tools in general and for specific scientific fields, as well as news on current developments and upcoming events in the RDM community.

Soon, *bwFDM* will expand its information services. A detailed mapping of RDM stakeholders, services and tools in Baden-Württemberg is underway. The mapping results will be provided on our future website, www.bwfdm.de (not yet available), which will be particularly useful for finding shared services.

Due to the growing demand for structured RDM training, we are designing an RDM certificate course tailored to specifics in Baden-Württemberg. The aim is to facilitate entry into RDM support at higher education and research institutions, whether in second-level support or in the role of a data steward.

Another *bwFDM*-flagship is the conference series *E-Science-Tage* with a strong interdisciplinary orientation. The conference series focuses on the topics of RDM and Open Science and provides diverse offerings for professional exchange between science and technology. Contributions connecting high performance computing with other aspects of research data infrastructure are most welcome for the next *E-Science-Tage* 2025 in Heidelberg.

Altogether, *bwFDM* supports the expansion of RDM activities in Baden-Württemberg, connecting various stakeholders and providing information and training.

Poster Session / 184

Base4NFDI - Basic Services for NFDI

Author: Franziska Fritzsche¹

¹ Gesis - Leibniz-Institute for Social Sciences

Corresponding Author: franziska.fritzsche@gesis.org

The aim of the German National Research Data Infrastructure (NFDI) is the systematic development of valuable data from science and research in order to link the decentralised and mostly inaccessible data with each other and to make them usable in a sustainable and quality-assured manner. In order to ensure a broad coverage of scientific disciplines, a total of 26 consortia from the cultural sciences, social sciences, humanities and engineering, as well as the life and natural sciences, have been selected in a science-led procedure steered by the German Research Foundation (DFG) over the past 3 years.

Since many scientific fields have similar requirements for research data and parallel developments

should be avoided as far as possible, all consortia are working together to support the Base4NFDI initiative. The goal is to develop, in cooperation with the community, basic services that enable sustainable research data management across disciplines. For this purpose a basic service is understood as a technical-organisational solution that usually includes storage and computing services, software, processes and workflows as well as the necessary personnel support for different service desks. The service should create added value for the consortia and their users, bundle existing services, be scalable, have a sustainable operating model, and be based on a standardised service model. Any idea for a basic service must originate from a community process within a section of the NFDI Association. If the proposal is successful, the development teams will receive support and guidance from the task areas of Base4NFDI in areas such as development, implementation and training. Once developed, the basic services will be made available to the scientific community. In this way, Base4NFDI actively contributes to the systematic opening and networking of the German science system.

Poster Session / 182

Further development of the state service "Scientific Data Storage"

Authors: Sabine Richling^{None}; Sven Siebler¹; Alexander Balz¹; Onno Nerjes¹; Florian Heuschkel^{None}; Martin Baumann¹

¹ Heidelberg University

Corresponding Authors: richling@urz.uni-heidelberg.de, martin.baumann@urz.uni-heidelberg.de

The state service (Landesdienst) Scientific Data Storage (SDS@hd) has been available for several years to researchers from institutions participating in the federated identity management in Baden-Württemberg (bwIDM). The service is used to store small and large-scale scientific data that is frequently accessed ("hot scientific data"). This poster gives an overview of the main service characteristics and presents the latest technical and organizational developments.

Poster Session / 188

bwForCluster Helix

Authors: Alexander Balz¹; Martin Baumann¹; Onno Nerjes¹; Sabine Richling¹; Sven Siebler¹

¹ Heidelberg University

Corresponding Author: martin.baumann@urz.uni-heidelberg.de

The bwForCluster Helix is a high performance computing system for researchers in Baden-Württemberg. It is designed to be used primarily in the fields of Structural and Systems Biology, Medical Science, Soft Matter, and Computational Humanities as well as for discipline-independent methods development. We present the architecture of the system and the interconnection with other services like SDS@hd and bwVisu.

Talks / 165

Introducing Galaxy: A new cloud and HPC gateway for data-driven computational research

Authors: Sanjay Kumar Srikakulam¹; Björn A Grüning¹

¹ *Albert-Ludwigs-Universität Freiburg*

Corresponding Authors: srikakus@informatik.uni-freiburg.de, gruening@informatik.uni-freiburg.de

Galaxy is a scientific workflow and data analysis platform transforming data-driven research. It is focused on creating an open infrastructure for computational research that is robust, scalable, and integrated, allowing for federated computational infrastructures and democratization of research data analysis. Galaxy's efficient web-based, intuitive user interface embedded with thousands of essential tools enables scientists to conduct sophisticated analyses without extensive programming knowledge or technical skills. By breaking down technical obstacles, Galaxy fosters innovation and collaboration across the scientific community. Galaxy excels in reproducibility, creating a research environment where findings can be validated and built upon and strictly FAIR. Extensive dataset collections, training materials, workflows, and expert-crafted tools enrich Galaxy, making it a valuable resource for researchers in various domains.

Galaxy is not just a platform for data analysis; it facilitates learning and collaboration through dedicated and unique Training Infrastructure as a Service (TaaS) and Virtual Research Environments (VREs). Through the EuroScienceGateway (ESG) project Galaxy serves as the Cloud and HPC gateway for computational research, empowering scientists with a user-friendly environment. Galaxy's versatility is evident in its application across diverse scientific fields, including life science, materials science, astrophysics, climate science, and more. In these domains, Galaxy aids in processing and analyzing complex problems, properties, and phenomena. This domain-agnostic nature highlights Galaxy's adaptability, contributing to its growing popularity among researchers from various disciplines.

With the European Galaxy server (<https://usegalaxy.eu>) serving 73K+ researchers globally, Galaxy stands as the new cloud and HPC gateway for domain-agnostic computational research, enabling scientists to navigate complex tools with ease without hindering their data analysis.

In this talk, we introduce Galaxy, its rich features and how Galaxy EU enables 73K+ researchers to conduct cutting-edge data-driven research by significantly reducing the computational challenges involved in complex data analysis.

Talks / 174

The agroecosystem multi-model library, Expert-N on high-performance computers (HPC)

Authors: Fasil Mequanint^{None}; Gayler Sebastian^{None}; Daniela Bendel^{None}; Thilo Streck^{None}

Corresponding Author: fasil.mequanint@uni-hohenheim.de

Soil-crop modeling plays a pivotal role in modern agricultural research by estimating environmental impacts, optimizing resource use, and predicting crop yields, particularly in the face of climate change.

Expert-N, an agro-ecosystem model library, provides a versatile framework for soil-plant-atmosphere modeling, leveraging High-Performance Computing (HPC) to enhance model accuracy through calibration and sensitivity analysis.

Here we describe the Expert-N software package and explore its integration with HPC for parallel computing, showcasing specific research endeavors undertaken in this domain.

Talks / 175

High-resolution global weather simulations with the Model for Prediction Across Scales (MPAS) on the Hawk HPC system

Author: Thomas Schwitalla^{None}

Co-authors: Volker Wulfmeyer ¹; Kirsten Warrach-Sagi ¹

¹ *Institute of Physics and Meteorology, University of Hohenheim*

Corresponding Author: thomas.schwitalla@uni-hohenheim.de

Recently, the frequency of droughts, heatwaves and other extreme events like heavy precipitation dramatically increased all over the world. Therefore, it is essential to accurately predict these types of events with a sufficient forecast lead time. As these extreme events are usually triggered by large-scale circulation pattern, this raises the question if the traditional high-resolution limited area model (LAM) approach to predict these events is still valid.

Current operational global numerical weather prediction (NWP) models operate on horizontal resolution in the range of 10 km. However, several recent studies have shown that a higher model resolution is required to increase the accuracy of heavy precipitation forecast and the prediction of clouds. The latter is especially important on the climate time scale as clouds have a major impact on the earths' radiation balance.

I will present first results of applying the NWP model MPAS on the global scale using horizontal resolutions of 3 km or less. In addition, I will give a brief introduction about implemented model enhancements, the I/O and the MPAS code performance on the HLRS Hawk system.