ORNL/SPR-2024/3620

Summary Report of the SOS26 Workshop held March 11-14, 2024

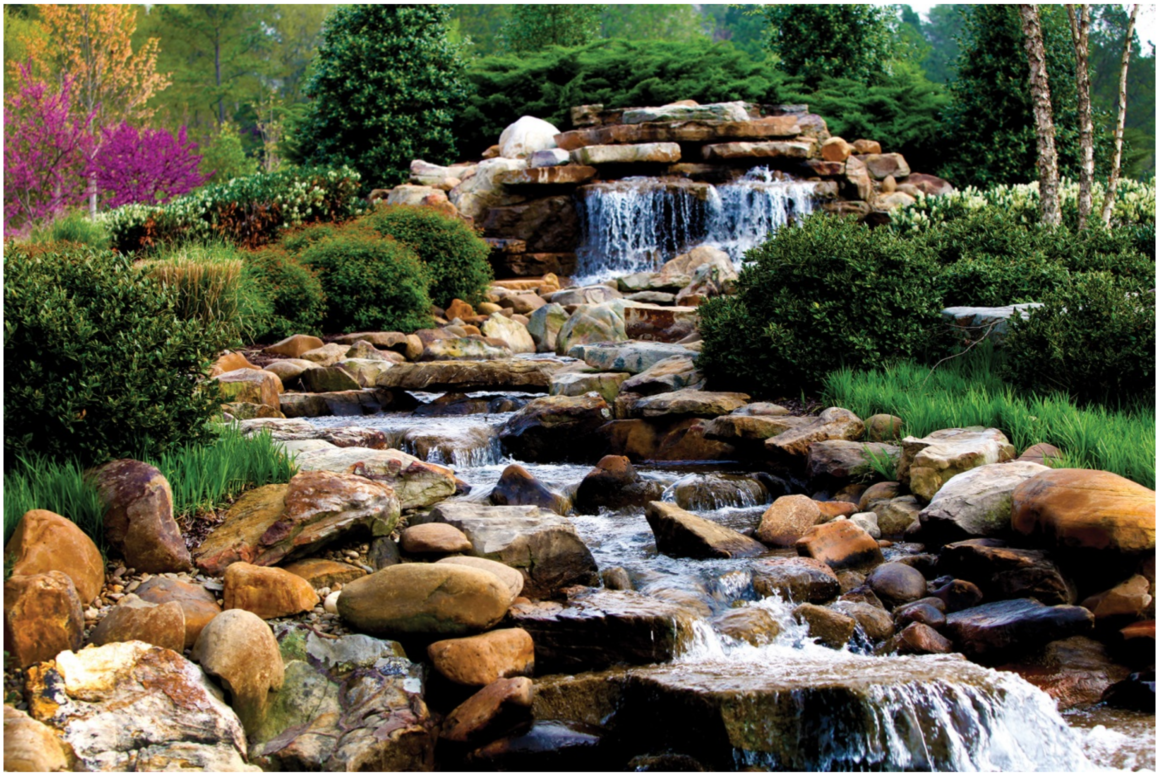

Ramanan Sankaran
Ron Brightwell
Joost VandeVondele
Gina Tourassi

September 2024

Oak Ridge National Laboratory

ORNL IS MANAGED BY UT-BATTELLE LLC FOR THE US DEPARTMENT OF ENERGY

ORNL/SPR-2024/3620

Computing and Computational Sciences Directorate

SUMMARY REPORT OF THE SOS26 WORKSHOP HELD MARCH 11-14, 2024

Ramanan Sankaran
Ron Brightwell
Joost VandeVondele
Gina Tourassi

September 2024

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-BATTELLE LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

Table of Contents

ORGANIZERS AND SPEAKERS

Organizers

Gina Tourassi, ORNL, Chair
Ron Brightwell, SNL, Co-Chair
Ramanan Sankaran, ORNL, Co-Chair
Joost VandeVondele, CSCS, Co-Chair

Session Chairs

Nageswara Rao, ORNL
Mike Parks, ORNL
David Bernholdt, ORNL
Juan Dorsch, CSCS
Veronica Melissa Vergara, ORNL
Kevin Pedretti, SNL
Jay Lofstead, SNL
Ron Oldfield, SNL
Stefano Schuppli, CSCS

Speakers

Casey Stone, ANL
Leonardo Sala, PSI/CH
Craig Bridges, ORNL
Larry Kaplan, HPE
Edoardo Paone, CSCS
Vadim Dyadechko, ExxonMobil
Nic Weber, U. of Washington
Hal Finkel, DOE/ASCR
Bjoern Enders, NERSC
Sean Wilkinson, ORNL
Kristian Ehlert, 4Cast
Andreas Fink, CSCS
Scott Atchley, ORNL
Debbie Bard, NERSC
Daniel Milroy, LLNL
Line Pouchard, BNL
Sarah Neuwirth, Uni Mainz
Michela Taufer, UTK

Katie Knight, ORNL
Rupak Biswas, NASA
Brian Ban Essen, LLNL
Philipe Ambrozi Dias, ORNL
Demitri Maestas, SNL
Ian Foster, ANL
Prasanna Balaprakash, ORNL
Venkatram Vishwanath, ANL
Steve Farrell, NERSC
Siva Rajamanickam, SNL
Jennifer Gaudioso, SNL

Logistical support

Alysha Tackett, ORNL
Lora Wolfe, ORNL

EXECUTIVE SUMMARY

The SOS26 workshop was organized by Oak Ridge National Laboratory (ORNL) and held March 11-14, 2024, at Cocoa Beach, Florida. The SOS is a workshop organized annually, with a focus on distributed high performance computing (HPC). The technical program of the workshop was developed jointly by Sandia National Laboratories (SNL), ORNL and the Swiss National Supercomputing Center (CSCS). The 2024 SOS26 workshop theme was "Versatile HPC for the evolving and expanding needs of science" and had seven technical sessions covering HPC, data and machine learning (ML) topics. Each session consisted of four or five presentations, followed by a panel discussion. This report documents the workshop proceedings covering all the technical sessions.

1. INTRODUCTION

1.1 Background and context

The SOS is a workshop organized annually with a focus on distributed high performance computing (HPC). The workshop series fosters a multi-lab and multi-national collaboration to explore the challenges faced in the development and use of large scale computing systems. The technical program of the workshop is developed jointly by Sandia National Laboratories (SNL), Oak Ridge National Laboratory (ORNL) and the Swiss National Supercomputing Center (CSCS). The workshop series began in 1997 through a recognition of the challenges faced by the HPC user facilities, vendors, computer scientists, system and application software developers, data and application scientists. The responsibility to host the workshop rotates among the three institutions (SNL, ORNL and CSCS) annually. The 2024 SOS26 workshop was organized by ORNL and was held March 11-14, 2024, at Cocoa Beach, Florida.

The origin of the workshop can be traced back to the concept of distributed supercomputing (DS) that developed in the 1990s by interconnecting commodity computing building blocks to create multi-Teraflop capability/capacity engines. The three organizing institutions saw DS as a complement to large-scale vendor machines and foresaw the development of a new computing industry. Based on this recognition, the first workshop in 1997 was centered on the theme of "Challenges and opportunities in connecting commodity components, potential from multiple vendors with commodity interconnects and software". Since then, the organizing committee has chosen a theme for the annual workshop based on developing HPC trends and community needs. The program is then organized to have multiple technical sessions with invited speakers from diverse backgrounds presenting a wide range of perspectives. At the conclusion of the workshop, a summary session is used to capture the proceedings and inform the theme of the next year's conference.

1.2 2024 Workshop format

The workshop theme in 2024 was "Versatile HPC for the evolving and expanding needs of science" and had seven technical sessions covering HPC, data and machine learning (ML) topics. The remainder of the report documents the workshop proceedings, and each session is included as its own section. The full agenda of the program is listed in Appendix A. Each session had one or two session chairs and four or five speakers. The technical sessions consisted of a panel discussion with audience participation following the individual presentations. The workshop also included a dinner banquet and an invited talk, which was sponsored by the sponsors acknowledged in Appendix B. A poster session was also organized during the workshop to encourage participation from early career researchers and students. At the conclusion of the workshop, the final session was used to summarize the workshop proceedings and provide the perspectives from the three organizing institutions.

2. SCIENCE INSTRUMENTS AND INTEGRATION

Instrument-computing ecosystems (ICE) have been proliferating rapidly and are finding application across many science disciplines, such as material science, chemistry, biology, etc. The proliferation of ICE has been powered in part by automation and artificial intelligence (AI). Integration of science instruments poses unique networking, software, workflow and orchestration challenges. Integration of instruments also has several network requirements consisting of physical network connectivity, routing, security zone alignment, firewalls and access layers for the applications and application program interfaces (API) used on the instruments. It is essential that the multiple networking layers be capable and aligned to allow communication between the software applications across instrument and system boundaries. There are several ICE projects at national labs and universities across the world. The workshop saw a selection of the projects present their experiences and perspectives.

2.1 Workcell Execution Interface (WEI)

WEI is a practical framework for the integration of diverse scientific instruments at scale for automated scientific experimentation. The integration of scientific instruments and robotics is often a barrier for physical scientists developing new autonomous discovery capabilities. The progress towards fully automated laboratories with the WEI was presented. WEI is an open, extensible and Python-centric software framework developed at Argonne National Laboratory as part of the Autonomous Discovery Initiative. WEI enables scientists to coordinate diverse instruments and services required for stepwise execution of various experimental workflows through the use of human-readable YAML formatted Workflow and Workcell definitions, as well as Python-based user interaction and logic implementation. Additionally, Docker containers are leveraged to facilitate the deployment and operation of instruments that can be controlled via various protocols, including REST, ROS2, and others, allowing for rapid integration of new instruments with minimal effort. WEI's ability to facilitate autonomous experimentation has been demonstrated. Efforts are underway to build out further closed-loop exemplar experiments in the fields of biology and material science.

2.2 Science Integration Challenges and Solutions at Paul Scherrer Institute (PSI)

PSI is facing challenges as a large-scale experimental research facility and presented their current experiences. PSI manages various kinds of experimental facilities, ranging from electron microscopes to brilliant X-ray sources. The constant advancement of accelerators, detectors and experimental methods translates into various information technology (IT) challenges ranging from high volume and throughput data storage to computational demands. Recent advances in ML and AI also pose very specific requirements on specialized hardware, such as requiring Graphic Processing Units (GPU) and Flexible Programmable Gate Arrays (FPGA). Various solutions are currently under evaluation at PSI to solve the integration challenges.

2.3 Autonomous Chemistry Laboratory (ACL)

Advanced research laboratories are producing ever more data and introducing more advanced methods of data analysis and feedback to control these laboratories. These "labs of the future" are intended to revolutionize the way that research is done, enabling in some cases the exploration of higher dimensional problems and better research outcomes. At ORNL, the Autonomous Chemistry Laboratory (ACL) is being developed as a self-driving laboratory to accelerate materials discovery and innovation. A combination of robotic tools is being integrated with computational tools to analyze reaction results and predict optimal directions for future investigation. This includes a mobile robot for flexibility in experimental configuration and sample transfers. The ability to combine aspects of liquid phase and solid-state synthesis enables a wide variety of synthetic experiments to be conducted, relevant to catalysis, energy storage, quantum materials, and more. The capabilities and concept for the ACL, along with challenges in building a flexible laboratory of this type, were presented.

2.4 Common Federation Framework

Domain sciences increasingly involve experiment steering in a federated and integrated workflow incorporating HPC simulation, AI training and inferencing, and Edge instrumentation. The Common

Federation Framework for Autonomous Instrumentation and Algorithmic Steering was presented. The various operational challenges were described and the approach of the proposed Common Federation Framework Workflow Software Development Kit (SDK) to enable this integration with federated diverse systems was presented. The presentation also included results showing quantitative improvements in a Scanning Transmission Electron Microscopy (STEM) spectrum reconstruction (Reptile-DKL) algorithm to be faster, more accurate, and generalizable while logging its metadata in an end-to-end manner.

3. SCIENTIFIC SOFTWARE SUSTAINABILITY AND STEWARDSHIP

Software is central in nearly every aspect of the computational science missions of agencies and institutions worldwide. Historically, questions of software quality and sustainability have mainly been left to the discretion of individual investigators. However, in recent years, there has been a growing recognition in the broader scientific software community that a more thoughtful approach to the stewardship of scientific software could be beneficial. The workshop had four presentations covering the perspectives from lab, industry, academia and sponsor agency.

3.1 Performance portability in weather and climate HPC applications

Productivity of domain scientists in running simulations on heterogeneous parallel architectures can be improved by decoupling application code from optimization logic. A framework for development and optimization of HPC applications for numerical weather prediction models was presented. It combines GT4Py, a Python Domain Specific Language (DSL) to efficiently express stencil computations, and DaCe, a data-centric approach to program optimization. By applying programmatic transformations on a dataflow graph intermediate representation of the program, the DaCe backend generates code tailored for the target architecture. This framework allows domain scientists to benefit from the highly productive Python ecosystem, while exploiting optimization opportunities unleashed by data-centric analysis.

3.2 Software sustainability in oil and gas industry

HPC is used extensively in the oil and gas industry, with primary applications being seismic imaging, reservoir modeling and simulation, and geomechanics/fracking simulations. Clearly, the HPC hardware has been getting faster, while the software running on such HPC systems gets much less attention and tends to stay largely in the shadow. Furthermore, the challenges in software sustainability remain subtle and are not well understood. There are several large questions facing software developers such as, "how can a software product survive the ever-changing computing landscape?" and "what would be a long-term winning strategy that can outlast short term benefits?". The people who develop, deploy and run the software dominate the productivity equation in HPC and are also a big risk factor. It would then seem that non-essential software complexity is to the detriment of sustainability and simplicity is the key to robustness. Post-release support and refactoring of software are unavoidable, consume a fair amount of resources and should be included in the plans for software development and sustenance.

3.3 The science of scientific software

Empirical studies are being used by the research community to learn about the science of scientific software, building on citation network studies. Recent results from quantitative studies based on bidirectional linking of repositories, funding awards, and journal articles were presented. Unique digital trace data from code repositories were used to describe patterns in the production, use, maintenance and abandonment of research software. Based on the analysis results, it was seen that about 25% of papers published on ArXiv link to version-controlled repositories, but only a small fraction of those repositories link back to the papers. Cases where there is a lack of plan for software sustenance by the developers were attributed to reasons such as "not ready for release", "no utility" and "lack of time". The historical pattern of software commits show that repositories that existed before an award date tended to have much longer periods of commit activity than repositories that were created during the award period.

3.4 Software stewardship and a next-generation software stack for scientific computing

The U.S. Department of Energy (DOE)'s Advanced Scientific Computing Research (ASCR) program has deployed the Nation's first exascale supercomputers and DOE is successfully completing the Exascale Computing Project (ECP). ECP created a portable state-of-the-art software stack for scientific high-performance computing that enables the portable targeting of a variety of current platforms – and future platforms, if the software stack continues to evolve as the technology ecosystem continues to advance. The presentation summarized ASCR's strategy for software stewardship, maintaining a vibrant community around its scientific-computing software ecosystem, as we look toward the future of scientific computing.

4. PROGRAMMATIC HPC ACCESS

Software is central in nearly every aspect of the computational science missions of agencies and institutions worldwide. Historically, questions of software quality and sustainability have mainly been left to the discretion of individual investigators. However, in recent years, there has been a growing recognition in the broader scientific software community that a more thoughtful approach to the stewardship of scientific software could be beneficial. The workshop had four presentations covering the perspectives from lab, industry, academia and sponsor agency.

4.1 The NERSC Superfacility API

The National Energy Research Scientific Computing (NERSC) center is a production HPC and data facility for the DOE Office of Science Research. The Superfacility concept is a framework for integrating experimental and observational instruments with computational and data facilities. Data produced by light sources, microscopes, telescopes, and other devices can stream in real-time to large computing facilities where it can be analyzed, archived, curated, combined with simulation data, and served to the science user community via powerful computing, storage, and networking systems. Connected with high-speed programmable networking, this superfacility model allows for discoveries across data sets, institutions, and domains and makes data from one-of-a-kind facilities and experiments broadly accessible. The superfacility API allows NERSC to embed HPC into cross facility workflows and provides a general purpose API for all NERSC users and projects to (i) plan/check availability of NERSC resource (ii) get data to NERSC when experiment is live (iii) start analysis job (iv) gather feedback quickly and (v) move data and results to archive after analysis.

4.2 Integrating OLCF HPC resources and scientific instruments with workflows

The Oak Ridge Leadership Computing Facility (OLCF) is a DOE Office of Science user facility offering leadership-class computing resources to researchers from government, academia and industry. The computational workflows for the present purpose might be defined as multi-step computational processes for working with data. It is recognized that the workflows are about automation, and the same elements that enhance human-machine interactions can also enable machine-machine interactions. The FAIR (findable, accessible, interoperable and reusable) guiding principles for scientific data when applied to workflows will have the attributes of both data and software. The presentation covered the current efforts towards applying FAIR at many levels of the ecosystem, namely, data, software, hardware, AI/ML and workflows. The presentation also covered OLCF's Integrated Research Infrastructure (IRI) efforts and test bed resources providing compute and storage. The IRI program is developing an ecosystem to support cross-facility workflows that help automate the hard parts – in secure, efficient, and performant ways – to radically accelerate discovery and innovation

4.3 Enhancing HPC service management using FirecREST API

As HPC evolves, there is an increasing need from the user community on creating and accessing sophisticated services on HPC. This talk presented the FirecREST API, a RESTful web API infrastructure that allows scientific communities to access the various integrated resources and services available on HPC systems. The RESTful API for HPC is needed by users that want to integrated complex services into HPC, need an abstraction layer for HPC resources and want to integrate with existing HPC services. FirecREST is a web enabled API to HPC resources, presenting a standard programming interface based on the RESTAPI concept. By staying independent of programming language, it translates web requests into HPC business logic while also parsing back HPC results into a web-friendly format. FirecREST was demonstrated using several use cases such as continuous integration (CI) pipelines, interactive computing, regression testing and workflow orchestration.

4.4 Mantik - A workflow tool for the development of AI on HPC

One of the challenges of today's ML workflows lies in the storage, exchange, and reproducibility of models. Furthermore, training ML models or utilizing them for inference on HPC systems is not standardized. The presentation introduced the Maelstrom project, which is providing the Mantik workflow tool that allows the development of and collaboration on ML models on HPC facilities like CSCS and Julich Supercomputing Center (JSC). In addition, the project supplies ML developers with a unified interface to HPC infrastructures, which enables training and deployment of ML models on different platforms. This tool integrates the existing open source framework, MLflow, to provide researchers with the ability to track their modeling results independent of any platform. Moreover, this tool enables storage of training workflows and models to ease sharing and ensures reproducibility.

5. ARTEFACT AND PIPELINE MANAGEMENT IN HPC

Enabling new types of science workflows on HPC platforms is an active area of research and there is significant interest in understanding the interaction of advanced workflows with containers, container orchestration, cloud and HPC integration, CI and continuous development (CD) pipelines, data management and digital thread management. End to end science requires enabling advanced workflow pipelines and artifact management on next-generation supercomputers. This session explored emerging advanced science workflows and how they are driving the technical requirements and architecture of future HPC platforms. The session was also aimed at gaining an understanding of the ongoing efforts, challenges and planned development in the community.

5.1 Scientific application deployment using CI/CD pipelines

CSCS aims to support different application lifecycles on the Alps HPC system, by providing the latest releases for software for users that need them, not forcing the users to update, while supporting the software through a full project life cycle of three to four years. Their application deployment using CI has three main components, namely (i) CI targeting HPC infrastructure (ii) user environments providing scientific applications to users and (iii) regression testing framework. The three components are combined via ReFrame to deploy containerized applications. ReFrame is a powerful framework for writing system regression tests and benchmarks, specifically targeted to HPC systems. The CI infrastructure at CSCS is deployed as microservices and middleware running on a Kubernetes cluster. The FirecREST API is used to submit and manage jobs running on the HPC systems. Finally, the OCI registry as storage (ORAS) system is used to manage the artifact outputs of the container build pipeline. The result is an end-to-end system for deploying reproducible, tested, and understandable HPC application environment.

5.2 OLCF's post-exascale system and thinking beyond Frontier

OLCF's fifth-generation system, Frontier, ushered in the exascale era in May 2022 and entered full production in early 2023. With the help of the Exascale Computing Project (ECP) and OLCF's Center for Accelerated Application Readiness (CAAR), Frontier's early science results have been impressive. The OLCF team, however, has already begun planning for Frontier's replacement. The presentation shared thoughts on OLCF's sixth-generation system, that follows Frontier. AI is growing in importance to OLCF's workload. In addition, DOE's ASCR office is pushing its computing facilities to more tightly integrate with DOE's experimental facilities to decrease the time needed to generate scientific insight. The presentation discussed these new drivers and their impact on the requirements for OLCF-6.

5.3 Complex workflows driving the architecture of the NERSC-10 system

One of the biggest challenges in scientific supercomputing today is the increasing complexity of scientific workflows and the new demands placed by them on traditional HPC infrastructure. The next generation of supercomputers will be designed to support workflows combining data movement and analysis, AI and large-scale simulations. The presentation described the complex HPC workflows that are driving the NERSC-10 system design. The presentation used case studies and real-life challenges to describe the science requirements that are driving the work, and how this translates to technical innovations for NERSC-10.

5.4 Enabling emerging scientific workflows with converged cloud and HPC environments

Emerging scientific workflows integrate software from traditional HPC, AI and ML, and supporting services. The complexity of the composite workflows translates to portability, deployment, and management challenges. Cloud technologies such as Kubernetes can facilitate the deployment and automate management of complex applications but were not designed for performance. The Flux Framework, a next-generation HPC resource management and scheduling framework developed at Lawrence Livermore National Laboratory (LLNL), uses hierarchical techniques to address challenges posed by complex scientific workflows and integrate cloud technologies. The presentation discussed research and development efforts on HPC, and cloud converged environments based on Flux and Kubernetes. The efforts seek to enable emerging scientific workflows by advancing converged computing: a best-of-both-worlds environment combining HPC performance with the automation, elasticity, and portability of the cloud.

6. REPRODUCIBILITY AND DATA MANAGEMENT

Enabling new types of science workflows on HPC platforms is an active area of research and there is significant interest in understanding the interaction of advanced workflows with containers, container orchestration, cloud and HPC integration, CI and continuous development (CD) pipelines, data management and digital thread management. End to end science requires enabling advanced workflow pipelines and artifact management on next-generation supercomputers. This session explored emerging advanced science workflows and how they are driving the technical requirements and architecture of future HPC platforms. The session was also aimed at gaining an understanding of the ongoing efforts, challenges and planned development in the community.

6.1 The role of FAIR in data-intensive, reproducible workflows

FAIR (Findable, Accessible, Interoperable, Re-usable) principles are underused by HPC and data-intensive communities who have been slow to adopt them. For machine learning applications, FAIR recommendations hardly exist. Numerous challenges to adopt FAIR in HPC include dataset sizes, lack of repositories, large and diverse metadata, complex heterogeneous execution environments, etc. A principled approach integrating FAIR with existing HPC tools and implementing at scale can help with such challenges. The presentation discussed provenance extraction and curation in a resource allocation system, metadata for performance analysis, and what reproducible AI could mean. These topics were illustrated in the context of RECUP, a framework for scalable metadata and provenance services for reproducible hybrid workflows. RECUP builds upon established HPC tools for scientific data management, proposes a set of tools to manage the entire life cycle of performance metadata, and aggregates them in an HPC-ready framework proposing concepts for reproducibility.

6.2 Parallel storage, monitoring tools and the quest for reproducibility

Parallel storage systems are sophisticated infrastructures designed to handle immense volumes of data concurrently. Their design is advancing beyond the traditional two-tiered distributed file system and archive model by introducing new tiers of temporary, fast storage close to the computing resources with distinctly different performance characteristics. Major challenges arise in maintaining coherence and minimizing latency in this intricate network, demanding advanced storage architectures, efficient data management, and synchronization techniques. Therefore, it is critical to bridge performance gaps through systematic monitoring, modeling, optimization, and analysis of the I/O and data paths in HPC systems to meet the colossal data demands of large-scale workloads. The presentation explored current parallel storage architectures and discussed challenges and opportunities in leveraging performance monitoring and modeling to support exascale I/O. The presentation also questioned to what extent reproducibility, replicability, comparability, and portability should influence the performance analysis cycle of parallel storage.

6.3 Democratizing data access and use through reproducibility, analytical tools and user services

The presentation introduced the National Scientific Data Fabric (NSDF), a cutting-edge platform designed to transform data management in combined HPC-cloud environments, aiming to democratize data access for enhanced scientific discovery. NSDF integrates user-centric networking, storage, and computing, simplifying HPC and cloud complexities for data analysts and boosting scientific reproducibility and clarity. The talk focused on the NSDF testbed, showcasing service integration for improved data management and scientific workflow reproducibility. It uniquely merges global and local networks, advanced storage, and powerful computing, ensuring interoperability and efficient data handling. A key part of the presentation was NSDF's application to the SOMOSPIE workflow, a machine learning model in earth science. This demonstrated the testbed's support for complex analyses in academic and commercial clouds. The SOMOSPIE case study within NSDF underscores its ability to handle complex data and its vital role in augmenting reproducibility and clarity in AI-driven research.

6.4 Reproducibility, reuse and other R words

The last R in FAIR (Findable, Accessible, Interoperable, and Reusable) is presented as one of the fundamental principles of digital data management for science. "Reuse" in this context means that data must be clearly described via domain-relevant metadata, retain traces of its provenance, and come accompanied by a clear usage license. This, in turn, can assist with experimental reproducibility by answering several questions. Where did the experimental evidence come from? If someone else is given access to the same data and applies the same methods, do they reach the same findings? However, what neither of these R words address is the problem of external validity. Given different reusable data, can you test your conclusions (replicate them) against other contexts? The presentation examined various R words and what the implications are for experimental data analysis, and why the HPC community should care.

7. LARGE LANGUAGE MODELS

Foundation models (FMs) and large language models (LLMs) are arguably the most disruptive technology of our generation. They have near-term impacts across all parts of our society including arts, industry, policy, cyber, intellectual property (IP) and law. More significant to the workshop and sponsors is the impact of these technologies on science, engineering and national security. The session consisted of speakers discussing topics ranging from large scale deep learning (DL), HPC training, AI strategies and model deployment.

7.1 FLASK: Foundation Learning AI for Synthesis Knowledge

FLASK is a new project that is focused on the development of a foundation model that can predict both novel molecules and novel pathways for synthesizing molecules. The current timeline for discovery and insertion of new materials is long and there is a need for tools that can accelerate this to meet DOE's design missions. Recent supply chain issues and emerging threats require discovering and manufacturing materials more rapidly. With a focus on the integration of graph-based representations with traditional transformer models, the question under test is if a dedicated synthetic chemistry model will outperform a more standard large-language model. The presentation outlined the trajectory of this new research project as well as some of the challenges for scaling on leadership-class HPC systems.

7.2 Multimodal foundation models for earth observation

Trained on vast unlabeled datasets through self-supervised learning, FMs enable generic feature extraction that facilitate specialization to a wide variety of downstream tasks. In addition to LLMs, FMs have proven effective in computer vision, as well as multimodal analytics pairing image and language. With the increasing availability of large Earth Observation (EO) data archives, there is a great opportunity to leverage FMs to aid real-world applications (e.g., disaster management). However, in contrast to most datasets used for FMs development and benchmarking, EO data imposes significant challenges related to the heterogeneity of data representations, sensor modalities, acquisition conditions, and overall geospatial concepts involved. The presentation discussed such challenges and opportunities regarding multimodal FMs for EO, the current state of the art on their development, and the latest progress/insights from the team at ORNL currently developing FMs for EO.

7.3 LLMs within the constraints of a secure, sand-boxed environment

There is a vision at Sandia National Laboratories (SNL) to bring LLM and natural language processing (NLP) capabilities as a service to SNL and its missions as quickly as possible, while carefully upholding software, data, and compute policies. The presentation focused on the capabilities and demonstrated use cases of applying emerging open-source NLP and LLM technologies to different applications within the constraints of a secure and sand-boxed environment while addressing LLM requirements and challenges such as compute, data access management, and hallucinations. These applications include Retrieval Augmented Generation (RAG), knowledge transfer, document clustering, and sentiment analysis. It also includes code generation, execution, translation, and interpretation.

7.4 AuroraGPT: A foundation model for science

The progression of scientific method has evolved to increasingly automated discovery processes, based on (i) large-scale mining of scientific knowledge at scale, (ii) AI assistants to generate hypotheses and propose protocols for testing those hypotheses, and (iii) autonomous computational and experimental systems to run those protocols. AI methods are important throughout these automated discovery processes. It is also hypothesized that many science use cases can be driven directly or indirectly from sufficiently powerful FMs. The presentation consisted of plans for AuroraGPT, a foundation model for science, and lessons learned in the project to date. The activities of the Trillion Parameter Consortium that has been established to support collaboration in this space was also discussed.

8. PLATFORMS FOR ML AT SCALE

The recent trend in ML has been to train ever larger language models. There are considerable discussions surrounding the unavailability of these models for further study. However, the hidden knowledge behind the efforts to train such large models has not gotten sufficient attention or discussion in prior workshops. The knowledge gained from executing a ML training workflow on thousands of GPU accelerated nodes is non-trivial and is worth sharing and documenting. This session covered the platforms and approaches for ML at scale and the benefit from scaling computational power to the highest levels.

8.1 Collaborative efforts in training large foundation models on Frontier

The first presentation was on the collaborative efforts of Oak Ridge National Lab, AMD, and Microsoft in training large AI foundation models for scientific research on the Frontier supercomputer. The AI foundation models in the discussion included: (i) HydraGNN, a multi-task learning graph neural network that utilizes extensive atomistic data for material discovery; (ii) a large-scale spatio-temporal transformer designed specifically for analyzing climate data, with the aim of enhancing the accuracy of climate forecasts; and (iii) a trillion-parameter language model developed to advance scientific research. This discussion underscored the significant impact of these advanced AI models in accelerating scientific discovery.

8.2 Scaling AI for science applications on novel architectures and systems

Improving the scaling of LLMs requires a two pronged approach with novel architectures as well as new algorithmic approaches. The experiences with scaling AI for science applications on the supercomputing systems at the Argonne Leadership Computing Facility (ALCF) were presented. Aurora is the flagship exascale HPC system at ALCF. AuroraGPT is an Argonne project with partners spanning industry and academia to pursue foundation models for open science and engineering. It is a series of dense models (7, 70, 200 and 1000 billion) trained on a mixture of general text, code and domain knowledge, including scientific papers and datasets with the aim of producing research grade artifacts for scientific community. The feasibility of training the AuroraGPT set of models on HPC systems was discussed. The presentation also covered the recent ALCF AI testbeds.

8.3 Enabling AI for science at scale on the Perlmutter supercomputer

The NERSC system roadmap for current Perlmutter system (NERSC-9) and future systems (NERSC-10, NERSC-11) has considered the user needs for AI training and inference, experiment data analysis along with HPC simulation and modeling. There is a trend of growing scientific AI workload on the NERSC systems with different types of models being trained using multiple frameworks as gleaned from the NERSC 2022 survey. Based on the current trends and future needs, the AI strategy is to (i) deploy optimized hardware and software systems, (ii) apply AI for science using cutting edge techniques and (iii) empower and develop workforce. The presentation discussed NERSC strategy for the pervasive AI ecosystem. Results from the MLPerf HPC benchmarks on Perlmutter were also presented.

8.4 Exploring LLMs for parallel code generation

Code generation for modern parallel architectures is a complex task that requires understanding of the architecture, parallel programming models, programming languages, and the context/domain. This is challenging for even human non-experts. The presentation discussed the use of LLMs for this task, especially using it for generating code in a co-pilot like environment. The focus of the presentation was mainly around the use of Kokkos programming model to generate parallel science codes that can be portable. Early efforts and initial results on this topic were presented.

9. SUMMARY

The presentations and discussions in this workshop identified several common themes and challenges that crosscut the HPC providers, developers and users across the globe. AI/ML and data intensive computing are posing additional requirements on the HPC providers for meeting the scientific user community needs. At the same time the frameworks and platforms for computing are also evolving rapidly providing opportunities for improved simulation and analysis workflows. Reproducibility, data management, tractable model training, instrument integration and security are some of the major concerns facing the community as we steadily progress in the exascale era.

APPENDIX A. AGENDA

APPENDIX A. AGENDA

Day 1 - Tuesday, March 12, 2024		
7:50 – 8:00 a.m.	Gina Tourassi, ORNL	Welcome

Session I: Science Instruments and Integration

Session Chairs: Nageswara Rao, ORNL, and Mike Parks, ORNL

8:00 – 8:10 a.m.	**Mike Parks, ORNL**	**Introduction**
8:10 – 8:30 a.m.	Casey Stone, ANL	Title: Workcell Execution Interface (WEI): A Practical Framework for Integration of Diverse Scientific Instruments at Scale for Automated Scientific Experimentation
8:30 – 8:50 a.m.	Leonardo Sala, PSI/CH	Title: Challenges for the next generation of experiments at large-scale research facilities
8:50 – 9:10 a.m.	Craig Bridges, ORNL	Title: Autonomous Chemistry Laboratory: Building an Integrated Laboratory for Liquid and Solid-State Synthesis
9:10 – 9:30 a.m.	Larry Kaplan, HPE	Title: Common Federation Framework for Autonomous Instrumentation and Algorithmic Steering
9:30 – 10:00 a.m.	Panel Discussion/Evaluation Documentation	

Session II: Scientific Software Sustainability and Stewardship

Session Chair: David Bernholdt, ORNL

10:30 – 10:40 a.m.	**David Bernholdt, ORNL**	**Introduction**
10:40 – 11:00 a.m.	Edoardo Paone, CSCS	Title: Performance portability in weather and climate HPC applications
11:00 – 11:20 a.m.	Vadim Dyadechko, ExxonMobil	Title: Outrunning the Red Queen
11:20 – 11:40 a.m.	Nic Weber, Univ. of WA	Title: Code Disposability: Recent Advances in the Science of Science Software

11:40 a.m. – 12:00 p.m.	Hal Finkel, DOE/ASCR	Title: Software Stewardship and a Next-Generation Software Stack for the Future of Scientific Computing
12:00 – 12:30 p.m.	Panel Discussion/Evaluation Documentation	
12:30 – 1:15 p.m.	Ramanan Sankaran, ORNL	Title: State of Scientific Computing and HPC

Session III: Programmatic HPC Access
Session Chair: Juan Dorsch, CSCS

1:30 – 1:40 p.m.	**Juan Dorsch, CSCS**	**Introduction to Session III**
1:40 – 2:00 p.m.	Bjoern Enders, NERSC	Title: The NERSC Superfacility API
2:00 – 2:20 p.m.	Sean Wilkinson, ORNL	Title: Integrating HPC Resources and Other Scientific Instruments with Workflows
2:20 – 2:40 p.m.	Juan Dorsch, CSCS	Title: Enhancing HPC Service Management using FirecREST API
2:40 – 3:00 p.m.	Kristian Ehlert, 4Cast	Title: Mantik – A workflow tool for the development of AI on HPC
3:00 – 3:30 p.m.	Panel Discussion/Evaluation Documentation	

3:30 – 5:00 p.m.	**Poster Session, Chair: Veronica Melesse Vergara, ORNL**

Day 2 - Wednesday, March 13, 2024		
7:50 – 8:00 a.m.	Ramanan Sankaran, ORNL	Overview of Day 2

Session IV: End-to-End Science: Enabling Advanced Workflow Pipelines and Artifact Management in Next-generation Supercomputers
Session Chair: Kevin Pedretti, SNL

8:00 – 8:10 a.m.	**Kevin Pedretti, SNL**	**Introduction to Session IV**
8:10 – 8:30 a.m.	Andreas Fink, CSCS	Title: Scientific application deployment using CI/CD pipelines
8:30 – 8:50 a.m.	Scott Atchley, ORNL	Title: Beyond Frontier – Thoughts on the OLCF's Post-Exascale System
8:50 – 9:10 a.m.	Debbie Bard, NERSC	Title: How complex HPC workflows are driving the architecture of the NERSC-10 system
9:10 – 9:30 a.m.	Daniel Milroy, LLNL	Title: Enabling Emerging Scientific Workflows with Converged Cloud and HPC Environments
9:30 – 10:00 a.m.	Panel Discussion/ Evaluation Documentation	

Session V: Reproducibility and Data Management
Session Chair: Jay Lofstead, SNL

10:30 – 10:40 a.m.	**Jay Lofstead, SNL**	**Introduction to Session V**
10:40 – 11:00 a.m.	Line Pouchard, BNL	Title: The role of FAIR in data-intensive, reproducible workflows
11:00 – 11:20 a.m.	Sarah Neuwirth, Uni Mainz	Title: Bridging the Gap: Parallel Storage, Monitoring Tools, and the Quest for Reproducibility
11:20 – 11:40 a.m.	Michela Taufer, UTK	Title: Advancing Science with the National Scientific Data Fabric: Democratizing Data Access and Use through Reproducibility, Analytical Tools, and User-Focused Services
11:40 a.m. – 12:00 p.m.	Katie Knight, ORNL	Title: Reproducibility, Reuse, and other R words

12:00 – 12:30 p.m.	Panel Discussion/ Evaluation Documentation

7:00 – 8:30 p.m.	**Invited Speaker:** **Rupak Biswas, NASA**	**Title of Presentation: HPC to Enable NASA Missions**

Day 3 - Thursday, March 14, 2024		
7:50 – 8:00 a.m.	Ramanan Sankaran, ORNL	Overview of Day 3

Session VI: Large Language Models

Session Chair: Ron Oldfield, SNL

8:00 – 8:10 a.m.	**Ron Oldfield, SNL**	**Introduction to Session VI**
8:10 – 8:30 a.m.	Brian Van Essen, LLNL	Title: FLASK: Foundation Learning AI for Synthesis Knowledge
8:30 – 8:50 a.m.	Philipe Ambrozi Dias, ORNL	Title: Challenges and opportunities of multimodal Foundation Models for Earth Observation
8:50 – 9:10 a.m.	Demitri Maestas, SNL	Title: Demonstrating the cutting-edge capabilities of Large Language Models (LLMs) within the constraints of a secure, sand-boxed environment
9:10 – 9:30 a.m.	Ian Foster, ANL	Title: AuroraGPT: A Foundation Model for Science
9:30 – 10:00 a.m.	Panel Discussion/ Evaluation Documentation	

Session VII: Platforms for ML at Scale

Session Chair: Stefano Schuppli, CSCS

10:30 – 10:40 a.m.	**Stefano Schuppli, CSCS**	**Introduction to Session VII**
10:40 – 10:55 a.m.	Mohamed Wahib, RIKEN	Title: Challenges of Scaling Deep Learning on HPC Systems
10:55 – 11:10 a.m.	Prasanna Balaprakash, ORNL	Title: Advancing Scientific Research with AI: Collaborative Efforts in

		Training Large Foundation Models on the Frontier Supercomputer
11:10 – 11:25 a.m.	Venkatram Vishwanath, ANL	Title: Scaling AI for Science Applications on Novel AI Architectures and Supercomputers
11:25 – 11:40 a.m.	Steve Farrell, NERSC	Title: Enabling AI for Science at Scale on the Perlmutter Supercomputer
11:40 – 11:55 a.m.	Siva Rajamanickam, SNL	Title: Exploring LLMs for parallel code generation
11:55 a.m. – 12:20 p.m.	Panel Discussion/ Evaluation Documentation	
12:30 – 1:00 p.m.	Ron Brightwell, SNL	Title: Enabling science using machine learning/AI

Session VIII: Crystal Ball and Invitation to SOS 27
Session Chair: Ramanan Sankaran, ORNL

1:00 – 1:10 p.m.	**Gina Tourassi, ORNL**	**View from Oak Ridge**
1:10 – 1:20 p.m.	Jennifer Gaudioso, SNL	View from Sandia
1:20 – 1:30 p.m.	Joost VandeVondele, CSCS	View from the Alps
1:30 – 1:45 p.m.	Wrap-Up Discussion/ Evaluation Documentation	
1:45 – 2:00 p.m.	Joost VandeVondele, CSCS	Announcement of SOS 27

APPENDIX B. SPONSORS

APPENDIX B. SPONSORS

(in alphabetical order)

1. Advanced Micro Devices (AMD)
2. ARM
3. Amazon Web Services (AWS)
4. CTG Federal
5. Cornelis Networks
6. DataDirect Networks (DDN)
7. Dell Technologies
8. General Electric (GE) Aerospace
9. Hewlett Packard (HP) Enterprise
10. International Business Machines (IBM)
11. Intel
12. nVidia
13. Penguin Solutions
14. Samsung
15. VAST