High-Performance Computing Highlights from the Department of Energy to be Presented at SC24

Researchers from across the U.S. Department of Energy (DOE) National Laboratory system will showcase their latest findings in computing at SC24, the International Conference for High-Performance Computing, Networking, Storage and Analysis taking place on Nov. 17–22 in Atlanta, GA.

The DOE has been a global leader in high-performance computing for decades, enabling novel discoveries in energy, materials sciences, artificial intelligence, and many other scientific disciplines. The DOE oversees the development of many of the world's most powerful supercomputers, which have been a critical component in multiple Gordon Bell Awards over the decades.

Attendees at SC24 can learn more about these awards and more at the DOE booth (#3401). Don't miss the featured talks, technical demonstrations, and roundtable discussions happening at the booth on Nov. 18–21. Topics include artificial intelligence, microelectronics, and energy-efficient computing.

Scheduled speakers and demos are noted below. Visit the DOE booth website (scdoe.info) for the latest information.

Featured Speaker Schedule:

Monday, Nov. 18

7: 45 p.m. **Thuc Hoang (National Nuclear Security Administration), Trish Damkroger (Hewlett Packard Enterprise), Brad McCredie (AMD), Rob Neely (Lawrence Livermore National Laboratory)**
"Introducing El Capitan: Exascale Computing for National Security"
Parts manufactured using laser-metal manufacturing techniques, such as laser welding and laser powder bed fusion additive manufacturing, are highly sensitive to the process details, which govern the formation of defects and distortions.

8:30 p.m. **Bronson Messer (Oak Ridge National Laboratory)**
"Fulfilling the promise of the world's first exascale supercomputer: Science on Frontier"

Tuesday, Nov. 19

10:45 a.m. **Luc Peterson (Lawrence Livermore National Laboratory)**
"Harnessing Exascale: Revolutionizing Scientific Discovery with El Capitan"
Imagine wielding tens of thousands of accelerated processing units (APUs – combined CPUs and GPUs with unified memory) in an architecture designed for both scientific computing and artificial intelligence—what frontiers could you explore? El Capitan, the first exascale machine developed by the Department of Energy's National Nuclear Security Administration (NNSA), stands on the brink of transforming our scientific and national security landscapes, boasting a staggering projected speed of over 2 exaflops.

11:30 a.m. **Malachi Schram (Jefferson Lab)**
"Advancing Science Discovery using AI at Jefferson Lab"

In this presentation, we will discuss some new efforts in data science at Jefferson Lab focusing on enabling science discovery, improving operations at DOE Scientific User Facilities (SUFs), and helping study regional issues. We will start by presenting a scalable asynchronous generative AI workflow designed for high-performance computing to solve fundamental nuclear physics questions which are at the heart of Jefferson Lab's mission. We will also explore how our research in AI/ML (e.g. conditional models, reinforcement learning, uncertainty quantification, etc.) can identify pending faults and improve operations at DOE SUFs. Finally, we will present our research in AI/ML that addresses important regional issues (e.g. flooding and health informatics).

1:00 p.m. Shantenu Jha (Princeton Plasma Physics Laboratory)
"HPC is dead. Long live HPC: AI-coupled High-Performance Computing for Design and Discovery"
It is well known that traditional HPC approaches for simulation & modeling are rapidly reaching limits. We will discuss why and how AI-coupled HPC is necessary to overcome these limitations. We will present multiple examples of AI-coupled HPC workflow for scientific discovery and design by researchers at PPPL and Princeton in a broad range of discipline domains.

1:45 p.m. Ian Foster (Argonne National Laboratory)
"AuroraGPT: Scaling AI for Scientific Discovery"
Argonne National Laboratory's AuroraGPT is an AI model designed to catalyze advancements in science and engineering that uses the Aurora exascale supercomputer. The primary goal is to build the infrastructure and expertise necessary to train, evaluate, and deploy large language models (LLMs) at scale for scientific research. Hear about what researchers have been doing to prepare the data used to train the model, efforts to evaluate the model, some early results, and how this project supports DOE's Frontiers in Artificial Intelligence for Science, Security and Technology initiative.

2:30 p.m. Shinjae Yoo, Yihui (Ray) Ren, Wei (Celia) Xu (Brookhaven National Laboratory)
"AI for Science: From Infrastructure to Scientific Application"
The rapid evolution of artificial intelligence (AI) is transforming scientific research across domains. This talk, featuring Brookhaven Lab AI Department researchers, explores the role of AI in revolutionizing science with an emphasis on both infrastructure and application. By addressing challenges from the hardware and infrastructure level to advanced applications, this talk will provide a comprehensive view of how AI can drive scientific discovery forward.

3:15 p.m. Earl Lawrence (Los Alamos National Laboratory)
"Artificial Intelligence for Mission"
The ArtIMis (Artificial Intelligence for Mission) Project is developing frontier AI to address mission challenges at Los Alamos National Laboratory. We are focused on two areas: (1) foundation models for science and (2) AI for design, discovery, and control. Methods are being developed and applied to four scientific domains: (a) materials performance and

discovery, (b) biosecurity, (c) applied energy, and (d) multi-physics systems. In addition, we're developing a cross-cutting capability in test, evaluation, and risk assessment for AI.

4:00 p.m. **Siva Rajamanickam (Sandia National Laboratories)**
"Turning the Titanic with a Leaf Blower: Co-designing programming models, libraries and advanced architectures"
In this talk, we will highlight some of our efforts for that past several years working with large and small hardware accelerator vendors in co-designing new architecture features, programming models, and software libraries. Our collaborations were at various levels from designing a backend in DOE performance portable software stack, new algorithms for linear algebra kernels, simulations to evaluate the usefulness of a new hardware feature for DOE applications. We will show examples from mini-applications, Gordon Bell finalists, production applications, and share some lessons learned for productive engagements in the future.

Wednesday, Nov. 20

10:45 a.m. **Woong Shin, Matthias Maiterth (Oak Ridge National Laboratory)**
"HPC Energy Efficiency @ OLCF"
In this presentation, we will explore the OLCF's journey from Summit to Frontier and beyond, focusing on enabling energy-efficient high-performance computing (HPC) for exascale achievements and the ongoing efforts towards the post-exascale era. The OLCF's dedication to energy efficiency is highlighted by its power-efficient building infrastructure and system architecture, which are continuously improved through the utilization of operational data. This process involves the adoption of AI/ML models, visual analytics, and digital twins, progressively evolving into a data and software-driven energy-efficient user facility.

11:30 a.m. **John Shalf (Lawrence Berkeley National Laboratory)**
"Energy Efficient Computing"
John Shalf will discuss the rise of specialized architectures to boost HPC performance in the post-Moore's Law era. Reminiscent of the so-called "Attack of the Killer Micros" heralding the arrival of microprocessors for HPC in the early 1990s, for the next generation of computing there is the potential for what could be called the "Attack of the Killer Chiplets" over the next five to 10 years to bring an era of HPC/AI modularity. This complements the emergence of similar in-memory or near-memory compute technologies such as Coarse-Grain Reconfigurable Arrays (CGRA), to enable specialized architectures, re-define computing and provide new avenues for advancing supercomputing speed and energy efficiency.

1:00 p.m. **Jordan Musser (National Energy Technology Laboratory)**
"Modeling and simulation of multiphase technologies"
This presentation provides an overview of the National Energy Technology Laboratory's (NETL) multiphase computational fluid dynamics codes. The highly successful Multiphase Flows with Interphase eXchanges (MFIX) suite has been used to model a wide range of

applications including post-combustion carbon capture, bioreactor optimization, and bio-FCC regeneration. MFIX-Exa, a state-of-the-art CFD code, developed under DOE's Exascale Computing Project, is built on the AMReX software framework and is designed to leverage modern accelerator-based compute architectures. This presentation further reviews the underlying physical models of both MFIX and MFIX-Exa and contrasts their similarities and differences.

1:45 p.m. **Aaron Fisher (Lawrence Livermore National Laboratory), Ramanan Sankaran (Oak Ridge National Laboratory), David Martin (Argonne National Laboratory)**
"HPC4EI: Bringing national lab scale supercomputing to US Industry"
DOE's High Performance Computing for Energy Innovation Program (HPC4EI) awards federal funding for public/private R&D projects aimed at solving key manufacturing challenges. Under the program, each selected industry partner gains access to the DOE National Labs' supercomputers and expertise to help them solve large scale problems with the potential for signification energy and CO2 emission savings. These collaborative projects help these industries become more competitive, boost productivity, and support American manufacturing jobs.

2:30 p.m. **Walid Arsalane (National Renewable Energy Laboratory)**
"HPC Multilevel User Environment Installation"
The National Renewable Energy Laboratory (NREL) has stood up the Lab's third generation HPC platform Kestrel. This talk will discuss how we provided a useful environment on Kestrel. In particular we will discuss difficulties and how they were mitigated. The environment was set up using a combination of methodologies including, integrating with the base Cray system, appropriate set up of module initiation, builds from source, binary installs and builds/installs with spack. The spack builds include a multilayer buildout which aides in deployment but is mostly hidden from the user.

3:15 p.m. **Graham Heyes (Jefferson Lab)**
"High Performance Data Facility Status and Plans"
In October of 2023, the DOE awarded the lead of the High-Performance Data Facility project to the Thomas Jefferson National Accelerator Facility in partnership with Lawrence Berkeley National Laboratory. This presentation will discuss where we are today and our plans for the future of the project.

4 p.m. **Ilya Baldin (Jefferson Lab), Yatish Kumar (ESnet)**
"Implementing Time-Critical Streaming Science Patterns Using Distributed Computational Facilities"
This talk highlights the achievements, the current state and the future directions of EJFAT (ESnet JLab FPGA-Accelerated Transport) project in enabling real-time seamless streaming of scientific data at Terabit rates between instruments and widely distributed processing facilities. This real-time linking of the experimental and computational resources of the U.S. research enterprise is central to the Department of Energy concepts of IRI (Integrated

Research Infrastructure) for the future of how science is done. It fully aligns with ESnet's vision for the role the network should play in supporting cutting edge research and is critical to the design of the High-Performance Data Facility (HPDF) led by JLab. In this talk we will describe the custom hardware developed by ESnet, as well as the software framework developed by ESnet and JLab to build this technologically unique solution.

Thursday, Nov. 21

10:45 a.m. **Antonino Tumeo (Pacific Northwest National Laboratory)**
"Bridging Python to Silicon: the SODA toolchain"
Systems performing scientific computing, data analysis, and machine learning tasks have a growing demand for application-specific accelerators that can provide high computational performance while meeting strict size and power requirements. However, the algorithms and applications that need to be accelerated are evolving at a rate that is incompatible with manual design processes based on hardware description languages. Agile hardware design tools based on compiler techniques can address these limitations by quickly producing an application-specific integrated circuit (ASIC) accelerator starting from a high-level algorithmic description. In this talk I will present the software-defined accelerator (SODA) synthesizer, a modular and open-source hardware compiler that provides automated end-to-end synthesis from high-level software frameworks to ASIC implementation, relying on multilevel representations to progressively lower and optimize the input code.

11:30 a.m. **Debbie Bard (National Energy Research Scientific Computing Center)**
"The DOE's Integrated Research Infrastructure (IRI) Program Overview"
This presentation provides an overview of the Department of Energy's Integrated Research Infrastructure (IRI) program, a comprehensive initiative designed to enhance and interconnect scientific research facilities across the nation. We will discuss how the IRI program aims to create a cohesive network of state-of-the-art laboratories, computational resources, and data sharing platforms. The talk will offer a brief overview of the program's mission and goals, and Technical Subcommittees (TS) work, with particular emphasis on the TRUSTID TS.

Technical Demonstrations

Monday, Nov. 18

7:00 p.m. **Demo Station 1:** Eli Dart (Lawrence Berkeley National Laboratory, Argonne National Laboratory)
"IRI Fusion Pathfinder Multi-Facility Demo"

Demo Station 2: Sutanay Choudhury (Pacific Northwest National Laboratory)

"AI-guided Hypothesis Generation and Design of Catalysts with Complex Morphologies and Reaction Networks"

8:00 p.m. **Demo Station 1:** Lois Curfman McInnes (Argonne National Laboratory)
"PESO: Partnering for Scientific Software Ecosystem Stewardship Opportunities"

Demo Station 2: Paul Lin (Lawrence Berkeley National Laboratory)
"Accurate in-situ in-transit analysis of particle diffusion for large-scale tokamak simulations"

Tuesday, Nov. 19

10:00 a.m. **Demo Station 1:** Lois Curfman McInnes (Argonne National Laboratory)
"PESO: Partnering for Scientific Software Ecosystem Stewardship Opportunities"

Demo Station 2: Anastasia Bernat, Tim Vega, Sai Munikoti (Pacific Northwest National Laboratory)
"Next-Generation AI Tools for Environmental Review and Permitting Efficiency"

11:00 a.m. **Demo Station 1:** Vardan Gyurjyan, Amitoj Singh (Jefferson Lab)
"Practical Hardware Accelerated Real-Time Multi-facility Streaming Workflow"

Demo Station 2: Joaquin Chung, Flavio Castro (Argonne National Laboratory)
"SciStream: Enabling Data Streaming between Science Instruments and HPC Nodes"

12:00 p.m. **Demo Station 1:** Sameer Shende (University of Oregon, Argonne National Laboratory)
"E4S: Extreme-scale Scientific Software Stack"

Demo Station 2: Valentine Anantharaj (Oak Ridge National Laboratory)
"A digital twin toward monitoring and forecasting of severe weather events in the future"

1:00 p.m. **Demo Station 1:** David Rogers (Oak Ridge National Laboratory)
"IRI Early Technologies and Applications Demos"

Demo Station 2: Derek Mariscal (Lawrence Livermore National Laboratory)
"Sidekick System: AI-enabled High Repetition Rate Laser experiment software"

2:00 p.m. **Demo Station 1:** Seongmin Kim (Oak Ridge National Laboratory)
"Distributed Quantum Approximate Optimization Algorithm for Large-Scale Optimization"

Demo Station 2: Benjamin Mintz (Oak Ridge National Laboratory)
"Interconnected Science Ecosystem (INTERSECT) for Autonomous Laboratories"

3:00 p.m. **Demo Station 1:** Sameer Shende (University of Oregon, Argonne National Laboratory)
"E4S: Extreme-scale Scientific Software Stack"

Demo Station 2: Paul Lin (Lawrence Berkeley National Laboratory)

"Accurate in-situ in-transit analysis of particle diffusion for large-scale tokamak simulations"

4:00 p.m. **Demo Station 1:** Amir Shehata (Oak Ridge National Laboratory)
"HPC/QC Integration Framework"

Demo Station 2: Alex Lovell-Troy (Los Alamos National Laboratory)
"OpenCHAMI: Open Source HPC System Management for future generations of DoE HPC"

5:00 p.m. **Demo Station 1:** Jean Luca Bez (Lawrence Berkeley National Laboratory)
"Drishti: I/O Insights for All"

Demo Station 2: Derek Mariscal (Lawrence Livermore National Laboratory)
"Sidekick System: AI-enabled High Repetition Rate Laser experiment software"

Wednesday, Nov. 20

10:00 a.m. **Demo Station 1:** David Rogers (Oak Ridge National Laboratory)
"IRI Early Technologies and Applications Demos"

Demo Station 2: Joaquin Chung, Caitao Zhan (Argonne National Laboratory)
"SeQUeNCe, a Customizable Discrete-Event Simulator of Quantum Networks"

11:00 a.m. **Demo Station 1:** Craig Vineyard, Christian Mayr (Sandia National Laboratories)
"The SpiNNaker2 Neuromorphic Computing Architecture – LLMs, Optimization, & AI/ML"

Demo Station 2: Imran Latif (Brookhaven National Laboratory)
"Advancing Sustainability in Data Centers: Evaluation of Hybrid Air/Liquid Cooling Schemes for IT Payload using Sea Water"

12:00 p.m. **Demo Station 1:** Craig Vineyard, Christian Mayr (Sandia National Laboratories)
"The SpiNNaker2 Neuromorphic Computing Architecture – LLMs, Optimization, & AI/ML"

Demo Station 2: Rohith Anand Varikoti (Pacific Northwest National Laboratory)
"CACTUS: Harnessing Open-Source LLMs and Domain-Specific Tools for Advanced Chemistry Reasoning"

1:00 p.m. **Demo Station 1:** Jim Brandt (Sandia National Laboratories)
"Data-Driven Autonomous Operations"

Demo Station 2: Ana Gainaru (Oak Ridge National Laboratory)
"I/O Innovations for Modern HPC Workflows"

2:00 p.m. **Demo Station 1:** Vardan Gyurjyan, Amitoj Singh (Jefferson Lab)
"Practical Hardware Accelerated Real-Time Multi-facility Streaming Workflow"

Demo Station 2: Ana Gainaru (Oak Ridge National Laboratory)
"I/O Innovations for Modern HPC Workflows"

3:00 p.m. **Demo Station 1:** Craig Vineyard (Sandia National Laboratories)
"TOPNMC: Exploring Neuromorphic Computing Impact"

Demo Station 2: Albert Vong (Argonne National Laboratory)
"Live Experiment-time Analysis of Advanced Photon Source Experiment Data using ALCF's Polaris Supercomputer"

4:00 p.m. **Demo Station 1:** Verónica G. Melesse Vergara (Oak Ridge National Laboratory)
"DOE-NIH-NSF Collaboration: Deploying Biomedical Retrieval Augmented Generation pipelines on Frontier as part of the NAIRR Secure Pilot"

Demo Station 2: Alex Lovell-Troy (Los Alamos National Laboratory)
"OpenCHAMI: Open Source HPC System Management for future generations of DoE HPC"

5:00 p.m. **Demo Station 1:** Sutanay Choudhury (Pacific Northwest National Laboratory)
"AI-guided Hypothesis Generation and Design of Catalysts with Complex Morphologies and Reaction Networks"

Demo Station 2: Valentine Anantharaj (Oak Ridge National Laboratory)
"A digital twin toward monitoring and forecasting of severe weather events in the future"

Thursday, Nov. 21

10:00 a.m. **Demo Station 1:** Anastasia Bernat, Tim Vega, Sai Munikoti (Pacific Northwest National Laboratory)
"Next-Generation AI Tools for Environmental Review and Permitting Efficiency"

Demo Station 2: Rohith Anand Varikoti (Pacific Northwest National Laboratory)
"CACTUS: Harnessing Open-Source LLMs and Domain-Specific Tools for Advanced Chemistry Reasoning"

11:00 a.m. **Demo Station 1:** Craig Vineyard (Sandia National Laboratories)
"TOPNMC: Exploring Neuromorphic Computing Impact"

Demo Station 2: Benjamin Mintz (Oak Ridge National Laboratory)
"Interconnected Science Ecosystem (INTERSECT) for Autonomous Laboratories"

12:00 p.m. **Demo Station 1:** Jim Brandt (Sandia National Laboratories)
"Data-Driven Autonomous Operations"

1:00 p.m. **Demo Station 2:** Max Lupo-Pasini (Oak Ridge National Laboratory)
"HydraGNN: a scalable graph neural network architecture for accelerated material discovery and design"

2:00 p.m. **Demo Station 1:** Amir Shehata (Oak Ridge National Laboratory)
"HPC/QC Integration Framework"

Demo Station 2: Max Lupo-Pasini (Oak Ridge National Laboratory)
"HydraGNN: a scalable graph neural network architecture for accelerated material discovery and design"

Learn more about the national lab activities at the DOE booth website. Find details about the SC24 conference — including a full schedule of events — at the SC24 website.