2024 OLCF USER MEETING

POSTER SESSION

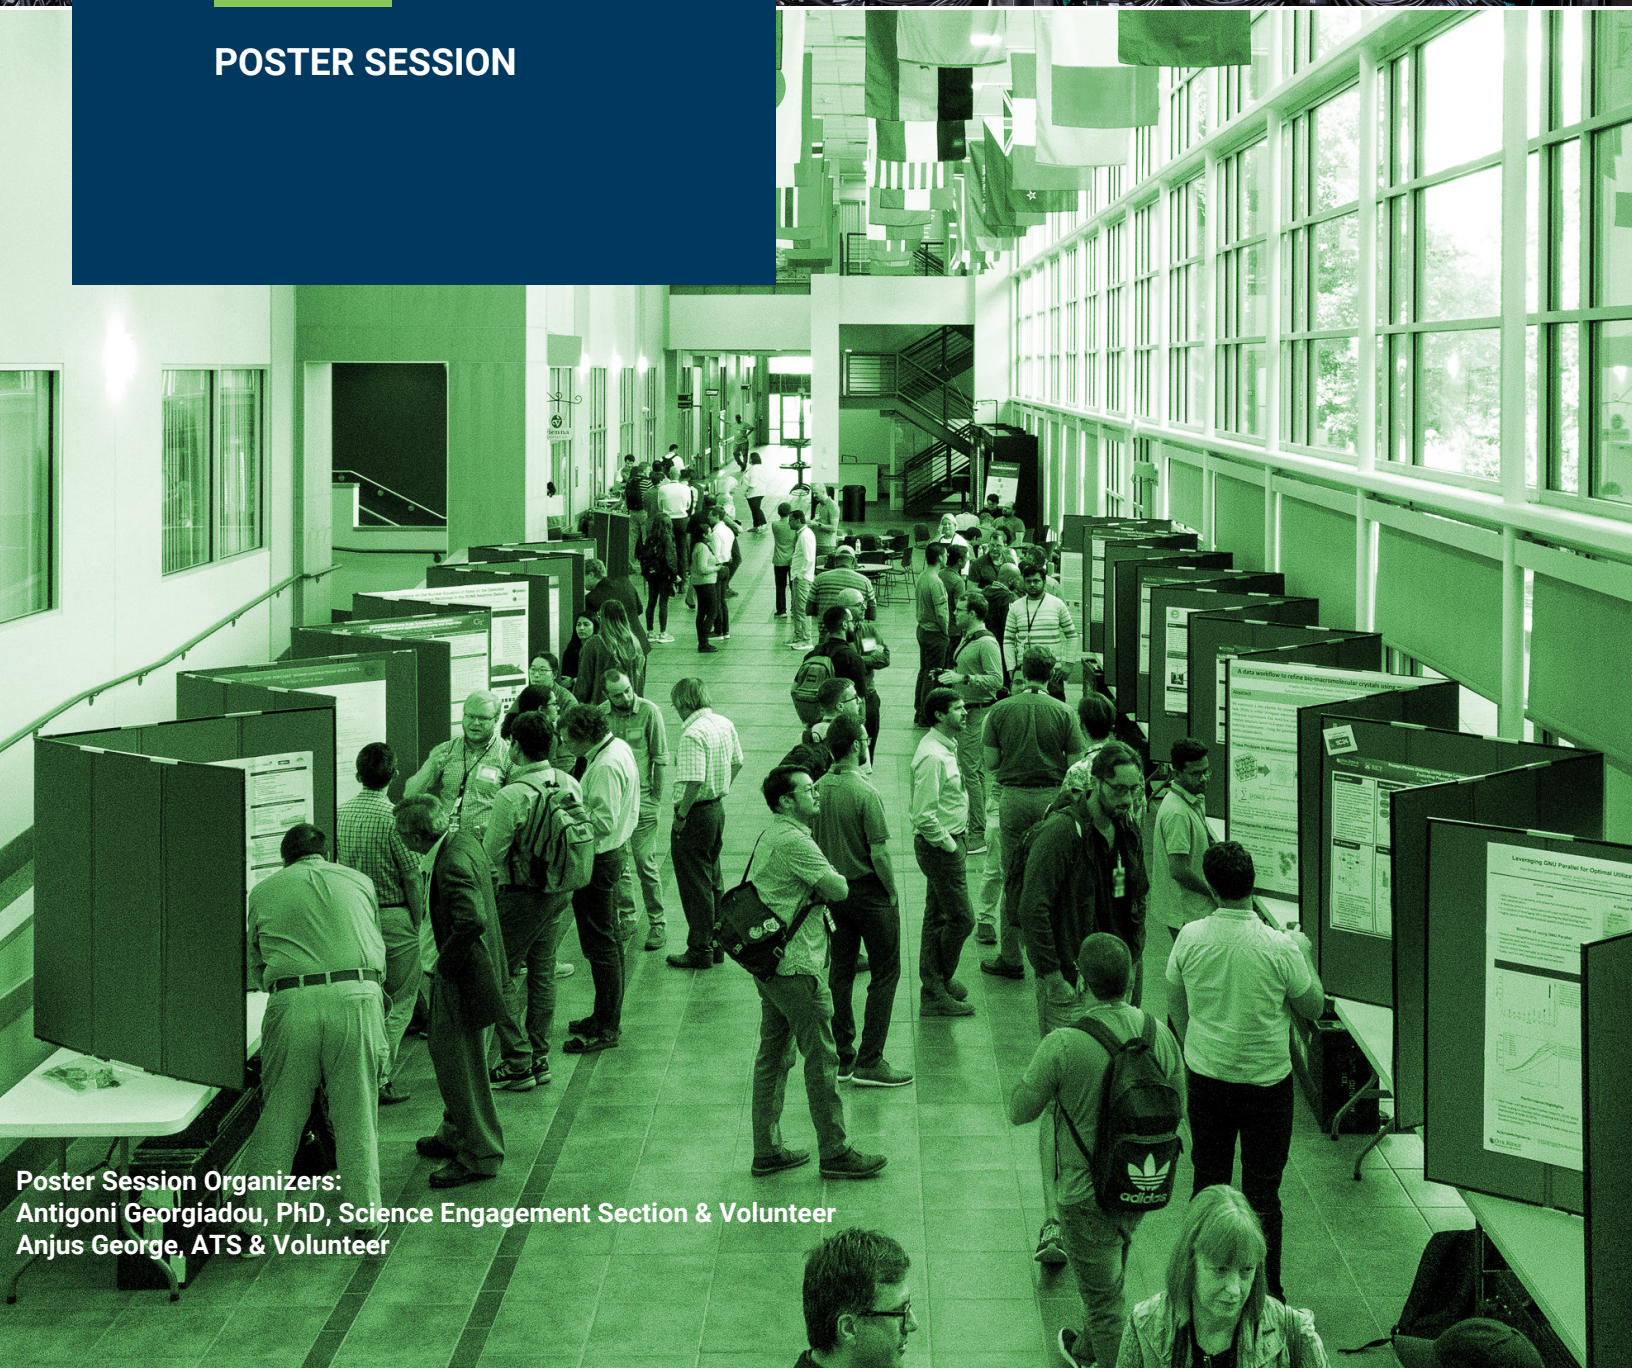

Poster Session Organizers:
Antigoni Georgiadou, PhD, Science Engagement Section & Volunteer
Anjus George, ATS & Volunteer

OAK RIDGE National Laboratory

2024-G04837/qjs

U.S. DEPARTMENT OF ENERGY

A Digital Twin Framework for Liquid-Cooled Supercomputers

PRESENTER
Matthias Maiterth

AUTHORS
Wes Brewer, Matthias Maiterth, Vineet Kumar, Rafal Wojda, Sedrick Bouknight, Jesse Hines, Woong Shin, Jake Webb, Scott Greenwood, Wes Wiliams, Daivd Grant, Feiyi Wang

ABSTRACT
“We have developed an open-source framework for developing comprehensive digital twins of liquid-cooled supercomputers, called “ExaDigiT”. It integrates three main modules: (1) a transient thermo-fluidic cooling mode, (2) a resource allocator and power simulator, and (3) and augmented reality model of the supercomputer and central energy plant. The framework enables the study of ``what-if” scenarios, system optimizations, and virtual prototyping of future systems. Using Frontier as a case study, we demonstrate the framework’s capabilities by replaying six months of system telemetry for systematic verification and validation and functional testing. ExaDigiT is able to elucidate complex transient dynamics of the cooling system. Run synthetic or real workloads, and predict energy losses due to rectification and voltage conversion. We envision the digital twin will be a key enabler for advancing sustainable and energy-efficient supercomputing.”

Exploring Julia as a unifying end-to-end workflow language for HPC on Frontier

PRESENTER
William F Godoy

AUTHORS
Caira Anderson, William F. Godoy, Pedro Valero-Lara, Katrina Lee, Ana Gainaru, Rafael Ferreira da Silva, Jeffrey S. Vetter

ABSTRACT
We evaluate Julia as a single language and ecosystem paradigm powered by LLVM to develop workflow components for high-performance computing. We run a Gray-Scott, 2-variable diffusion-reaction application using a memory-bound, 7-point stencil kernel on Frontier, the US Department of Energy’s first exascale supercomputer. We evaluate the performance, scaling, and trade-offs of (i) the computational kernel on AMD’s MI250x GPUs, (ii) weak scaling up to 4,096 MPI processes/GPUs or 512 nodes, (iii) parallel I/O writes using the ADIOS2 library bindings, and (iv) Jupyter Notebooks for interactive analysis. Results suggest that although Julia generates a reasonable LLVM-IR, a nearly 50% performance difference exists vs. native AMD HIP stencil codes when running on the GPUs. As expected, we observed near-zero overhead when using MPI and parallel I/O bindings for system-wide installed implementations. Consequently, Julia emerges as a compelling high-performance and high-productivity workflow composition language, as measured on the fastest supercomputer in the world. The showcased work was recognized with a best paper award at the SC23 WORKS workshop.

Performance Profiling of the GronOR Non-Orthogonal Configuration Interaction Application

PRESENTER
Xinju (Spancer) Dong

AUTHORS
Xinju (Spancer) Dong, Tjerk Straatsma, Coen de Graaf, Ria Broer, Carmen Sousa, Jordi Ribas, Xavier López

ABSTRACT
The popularity of non-orthogonal wavefunction methods keeps growing recently due to their potential of capturing correlations in a much more intuitive way compared to orthogonal approaches. Non-orthogonal configuration interaction using the fragment wave functions (NOCI-F) is one of such methods constructed for providing a novel computational perspective on processes such as designing new catalytic, electronic or photovoltaic materials. The GronOR software is built for this purpose. However, the high computational cost of using non-orthogonal wavefunctions is an unavoidable challenge beyond the small chemical systems. By taking advantage of modern HPC systems, GronOR relies on an advanced worker-master working model, which maximizes the efficiency by computing the Hamiltonian matrix in parallel. This execution model in GronOR has been applied successfully on multiple HPC systems, including Summit and Frontier. In our study, singlet fission process of indolonaphthyridine dimer is explored using GronOR on Summit and Frontier to calculate the electronic couplings between states involved in crucial electron excitations in potential photovoltaic materials. Focus of the current work is on the performance of the GronOR software on Summit.

Score-PDB: A diffusion model to predict crystallographic phases

PRESENTER
Mijanur Palash

AUTHORS
Mijanur Palash, Khadiza Begam, Zach Fox, Junqi Yin, and Jens Glaser

ABSTRACT
In this work, we present Score-PDB, a score-based diffusion model for crystallographic applications. By leveraging the gradient of the data distribution as a score function, Score-PDB can model diffraction data accurately. It denoises under-sampled data, regenerates missing values and improves image quality. Moreover, using experimental data as a guide, Score-PDB reconstructs missing phase information using only amplitude data, which is a crucial step in crystal structure factor refinement. In preliminary experiments on Frontier with 512 nodes and the RCSB Protein Data Bank (PDB) dataset, our model shows promising results and scaling capacity.

A data workflow to refine bio-macromolecular crystals using machine learning

PRESENTER
Khadiza Begam

AUTHORS
Mijanur Palash, Zach Fox, Junqi Yin, and Jens Glaser

ABSTRACT
We implement a data pipeline for training a novel crystallographic phasing model which leverages machine learning on the protein data bank (PDB) to refine biological macromolecular structures. Phasing has been a longstanding problem, complicating the interpretation of diffraction experiments. Our workflow uses the crystallographic toolbox (CCTBX) to generate electron density maps and corresponding complex structure factors in k-space from protein crystal structures, which encode information about atomic positions in X-ray and neutron scattering experiments. Using the generated data, we train a diffusion model, Score-PDB, which will have applications in drug discovery and bio-preparedness.

Harnessing Contrastive Learning for Drug Discovery

PRESENTER
Maria Batool

AUTHORS
Maria Batool, Verónica Melesse Vergara, and Jens Glaser

ABSTRACT
Contrastive learning, a self-supervised machine learning technique, offers a promising approach to uncover drug potential by modeling the correlation between molecular sequences and biological activity. The Contrastive Language-Image Pre-training (CLIP) method for images and text captions has demonstrated that learned representations are highly transferable to generative tasks when trained on multi-modal data. In drug discovery, similar large, multi-modal datasets comprising sequences or chemical structures of molecules are available. We use datasets from different online databases such as PDBBind, CHEMBL, and Binding DB containing up to 5M drugs and targets, along with their corresponding affinity values, i.e., dissociation or inhibitory constants. We obtain feature representations from both drug (SMILES string) and target (amino acid sequence) molecules as inputs to a transformer encoder. Our results show that contrastive fine-tuning offers an efficient way to induce semantic structure into the latent space of small molecules and proteins for downstream drug discovery tasks such as regression. We test the generalization ability of the method across diverse datasets.

MeshFormer: Hybrid Graph-Transformer for Learning Physical Simulations in Unstructured Meshes and Varying Geometries

PRESENTER
Sheikh Md Shakeel Hassan

AUTHORS
Hassan, Sheikh Md Shakeel and Gounley, John and Laiu, M. Paul and Zhang, Pei

ABSTRACT
Modern data-driven surrogates for time-dependent PDEs provide a fast alternative to expensive physics-based simulations often achieving speedups of up to 1-2 orders. Vision Transformers have proven to be effective in modeling such spatio-temporal physical systems. But, they fail to handle arbitrary geometries and domains that are often encountered in mesh-based simulations. We develop a hybrid Graph Neural Network and Transformer architecture that can handle arbitrary geometries and domains. Message Passing GNNs are utilised to extract rich spatial features from input meshes which are then passed into a Transformer network that further learns spatial and temporal relationships in the high dimensional embedding space. The model is evaluated for two CFD datasets, cylinder flow and flow over airfoils.

Advanced High-Performance Computing Techniques for Optimizing Cisplatin Delivery: Developing Predictive Models for Membrane Permeability and Radiosensitization

PRESENTER
Maria Batool

AUTHORS
Maria Batool, Van Ngo, John Vant, Anuj Kapadia, Debsindhu Bhowmik

ABSTRACT
“Cisplatin, a platinum-based chemotherapy drug, remains a fundamental treatment for cancer; however, its effectiveness is often compromised by resistance mechanisms and inadequate permeability through cancer cell membranes. This study explores how high-performance computing (HPC) can enhance cisplatin’s efficacy by investigating its transport through asymmetric cancer lipid membranes and its interactions with radiosensitizers. We employ HPC to model the molecular dynamics of cisplatin transport, aiming to devise strategies that improve its delivery to cancer cells. Specifically, our approach involves analyzing the lipid compositions of cancer cell membranes, with a focus on asymmetrical distributions that hinder drug permeability. By leveraging the computational power of ORNL’s exascale computing facilities, we conduct large-scale molecular dynamics simulations, providing atomic-resolution insights into how cisplatin and radiosensitizers interact with these membranes. Our use of HPC enables detailed simulations of complex membrane environments, revealing energy barriers and diffusion processes that affect drug transport. We build a comprehensive library of ligand permeability data, which, coupled with advanced machine learning algorithms, helps us develop predictive models for membrane permeability and optimize drug formulations. The integration of HPC with experimental validation through in vitro models allows us to test and refine these predictions. This research has the potential to revolutionize treatment paradigms

by enhancing drug delivery, overcoming resistance mechanisms, and ultimately improving patient outcomes. By combining computational modeling with experimental techniques, we aim to advance both chemotherapy and radiotherapy strategies, offering new insight”

The Kokkos Ecosystem

PRESENTERS
Bruno Turcksin, Damien Lebrun-Grandie, Daniel Arndt, Andrey Prokopenko

AUTHORS
Bruno Turcksin, Damien Lebrun-Grandie, Daniel Arndt, Andrey Prokopenko

ABSTRACT
The Kokkos ecosystem is a C++ performance portable ecosystem designed to bridge the gap between complex algorithms and diverse hardware architectures. It provides a unified programming model that abstracts away the intricacies of different computing platforms, allowing to write code once and run it efficiently on both CPUs and GPUs. At its core, Kokkos offers abstractions for parallel execution and memory management, enabling developers to focus on problem-solving rather than low-level hardware details. Its flexibility stems from its ability to target various backend execution models, including CUDA, HIP, SYCL, OpenMP, OpenACC, etc.

Visualizing Power Outages using EAGLE-I Data

PRESENTER
Edwin Washington

AUTHORS
Edwin Washington

ABSTRACT
In May 2024, a series of tornadoes struck Tallahassee, FL, causing widespread damage and power outages, significantly impacting the local community, including Florida A&M University. This study leverages EAGLE-I™ National Outage Map data, which monitors public electric utility customer outage websites and updates every 15 minutes, to analyze the extent and duration of power outages at the zip code and county levels. By integrating weather data with outage reports, we aimed to improve the accuracy of outage duration estimates and understand the correlation between severe weather events and power disruption. The analysis focused on Leon County, where significant outages were observed, with over 10% of customers losing power in several instances. Visualizations created using Python’s Pandas library highlighted the relationship between the storm’s path and the resulting outages. Future work includes refining the data models at the zip code level and expanding the analysis to other weather-related events. This research contributes to enhancing EAGLE-I’s capability to provide more localized and accurate outage data, ultimately aiding in more effective emergency response and utility management.

Prompt Phrase Ordering Using Large Language Models in HPC: Evaluating Prompt Sensitivity

PRESENTERS
Noah Thomasson and Hilda B. Klasky

AUTHORS
Noah Thomasson and Hilda B. Klasky

ABSTRACT
Large language models (LLMs) often require well-designed prompts for effective responses, but optimizing prompts is challenging due to prompt sensitivity, where small changes can cause significant performance variations. This study evaluates prompt performance across all permutations of independent phrases to investigate prompt sensitivity and robustness. We used two datasets: GSM8k, for mathematical reasoning, and a custom prompt for summarizing database metadata. Performance was assessed using the llama3-instruct-7B model on Ollama and parallelized in a high-performance computing environment. We compared phrase indices in the best and worst prompts and used Hamming distance to measure performance changes between phrase orderings. Results show that prompt phrase ordering significantly affects LLM performance, with Hamming distance indicating that changes can dramatically alter scores, often by chance. This supports existing findings on prompt sensitivity. Our study highlights the challenges in prompt optimization, indicating that modifying phrases in a successful prompt does not guarantee another successful prompt.

Examining Epidemic Models and Public Health Policy Effectiveness

PRESENTERS
Gabriella Harrison and Violet YarKhan

AUTHORS
Gabriella Harrison and Violet YarKhan

ABSTRACT
As a part of the NGP Internship program, an exploration of epidemic models was conducted. During the investigation, agent-based Monte-Carlo techniques and SIR Model techniques were compared to determine the most realistic model for demonstrating the effectiveness of various public health policies. The population was split into blue collar, white collar, and stay-at-home working classes to provide more detailed data. Then, using a combination of the two modeling techniques, it was shown that social distancing policies are most effective in the short-term and vaccination policies are most effective in the long-term. Overall, the ideal combination of health policies involves both social distancing and vaccination policies. This project provides insight into disease modeling in the modern world, and the impacts it can have on policy-making.

An Autoencoder-based Lossy Compression Method for E3SM Data

PRESENTER
Alexander Yan

AUTHORS
Alexander Yan, Qian Gong, Viktor Reshniak

ABSTRACT
“The Exascale Earth System Model (E3SM) simulates Earth’s climate over long timescales (20-100 years), producing vast amounts of data (10-100 TB) that require storage and sharing. Machine learning (ML) techniques are increasingly applied in climate science for tasks like data compression, feature tracking, and numerical weather forecasting. However, ML models depend on large datasets to improve accuracy, making storage and sharing of these datasets critical. Storing such large volumes on high-speed filesystems is costly, and transferring them across geographically distributed facilities is time-consuming due to limited network speeds. Data compression is a viable solution to these challenges, with lossy compression becoming more popular for its ability to reduce floating-point scientific data, which are difficult to compress losslessly. Numerical lossy compressors like ISABELA, SZ, ZFP, and MGARD can significantly reduce data while keeping compression-induced errors within acceptable bounds. ML-based compression approaches are also gaining traction for their nonlinear parameterization and ability to retain task-specific information. However, when designing a lossy compression algorithm, it is essential to carefully assess the errors and their impact on derived quantities to ensure the reliability of scientific results. In this work, we developed an AutoEncoder-based lossy compression model tailored for E3SM data. We evaluated its compressibility and used the feature tracking package TempestExtreme (TE) to examine the impact of lossy compression on post-analysis, informing the design of a trustworthy ML-based lossy compression model that balances data reduction with scientific integrity.”

Molecular Modeling of Antifouling Polymers for Biological Applications

PRESENTER
Seonghan Kim

AUTHORS
Seonghan Kim, Zhefei Yang, Jacek Jakowski, Jan-Michael Carrillo, Panchapakesan Ganesh, and Scott Retterer

ABSTRACT
Antifouling polymers are highly valuable in various applications, including biofouling-resistant coatings, medical devices, and nanoscale therapeutic/diagnostic materials. In nanomedicine, proteins in biological environments can interact with the surface of nanoparticles, altering their properties by changing its chemical environment and leading to the formation of protein corona structures. Preventing this unwanted interaction can be achieved by creating antifouling surfaces. While extensive experimental studies have characterized the surface interactions with proteins, computational insights remain limited due to challenges in modeling these structures. In this study, we present all-atom model structures of antifouling polymers grafted onto quartz surfaces and compare their interactions with proteins through molecular dynamics simulations.

Sub-grid parameter calibration for Cosmological simulations

PRESENTER
Antigoni Georgiadou

AUTHORS
James Ahrens, JD Emberson, Nick Frontiere, Antigoni Georgiadou, Salman Habib, Katrin Heitmann, David Higdon, Patricia Larsen, Chloe Keilers, Earl Lawrence, Zarija Lukić, Sandeep Madireddy, Kelly Moran, Nesar Ramachandra, Jean Sexton

ABSTRACT
Upcoming wide-field cosmic surveys are extremely promising to help in addressing the major challenges of cosmology, e.g. understanding the nature of the dark matter constituent of the universe, its accelerated expansion and model the dynamics of the universe. The strength in depth and width of these surveys will result in an unprecedented flood of data of large-scale structure, and theoretical models should be prepared to meet such a precision in astronomical information. At the same time cosmological simulations can potentially produce data in the range of several petabytes of data in a single run. As exascale computing becomes predominant, simulations from HACC are expected to produce 10 to 100 times as much data. Thus the team works alongside computer manufacturers and software developers to prepare as efficient and coherent exascale platforms as possible. The HACC code will produce science to tackle challenging cosmological problems by producing very large cosmological simulations run with the HACC cosmology code to address multiple-level science problems of interest. To achieve that a series of various sizes simulations will be created to test initial value problems that best suit the scientific goals.In particular the HACC code is using the INCITE program and Frontier to produce cosmological simulations of different implementations: (1) large-volume, high-mass gravity-only simulations, (2) large-volume, high-mass hydrodynamic simulations, (3) small-volume, very high-mass hydrodynamic simulations.

Solving the Hele-Shaw flow using the Harrow-Hassidim-Lloyd algorithm on superconducting devices: A study of efficiency and challenges

PRESENTERS
Eduardo Antonio Coello Pérez, Kalyan Gottiparthi

AUTHORS
Muralikrishnan Gopalakrishnan Meena, Kalyan Gottiparthi, Justin Lietz, Antigoni Georgiadou, and Eduardo Antonio Coello Pérez

ABSTRACT
The development of quantum processors capable of handling practical fluid flow problems represents a distant yet promising frontier. Recent strides in quantum algorithms, particularly linear solvers, have illuminated the path toward quantum solutions for classical fluid flow solvers. However, assessing the capability of these quantum linear systems algorithms (QLSAs) in solving ideal flow equations on real hardware is crucial for their future development in practical fluid flow applications. In this study, we examine the capability of a canonical QLSA, the Harrow-Hassidim-Lloyd (HHL) algorithm, in accurately solving the system of linear equations governing an idealized fluid

flow problem, specifically the Hele-Shaw flow. Our investigation focuses on analyzing the accuracy and computational cost of the HHL solver. To gauge the stability and convergence of the solver, we conduct shots-based simulations on quantum simulators. Furthermore, we share insights gained from executing the HHL solver on superconducting quantum devices. To mitigate errors arising from qubit measurement, gate operations, and qubit decoherence inherent in quantum devices, we employ various error suppression and mitigation techniques. Our preliminary assessments serve as a foundational step towards enabling more complex quantum utility scale evaluation of using QLSA for solving fluid flow problems.

Supporting Growth in the Quantum Computing Community

PRESENTER
Suzanne Prentice

AUTHORS
Suzanne Prentice, Ryan Landfield, Michael Sandoval

ABSTRACT
Supporting growth in Quantum Computing is a necessary, but challenging task to accomplish; the Oak Ridge Leadership Computing Quantum Computing User Program (OLCF QCUP) uses a clear approach that is simple both for users to understand and for other facilities to model. Firstly, the program aims to enable research by providing a broad spectrum of user access to the best available quantum computing systems, demonstrated by the fact it has over 370 users working on more than 130 projects, exploring a wide array of fields. Users need no prior experience and can choose to request access to a variety of machines such as numerical simulators and general-purpose gate systems. The QCUP streamlines access, providing technical and scientific support similar existing high performance computing approaches. This makes it easier for existing HPC users to onboard and explore quantum computing projects. Secondly, outside of direct user and research support it is necessary for growth to evaluate current quantum technologies and use cases. QCUP both leads integration of quantum computing with existing scientific workflows and systems, and it monitors metrics and benchmarks that track progress in system utility to evaluate the performance of early quantum computing applications. Lastly, it is essential to support growth of the quantum ecosystem by engaging with users, developers, vendors, and providers. QCUP accomplishes this through engagement with multiple populations of end users and stakeholders; it leads interactions between users, developers, vendors, and providers through hackathons, trainings, and forums to bring the community together. By sharing this information and best practices across user communities, the computing community is made more resilient.

Quantum simulations of Many-body Neutrino Fast-Flavor conversions in Supernovae.

PRESENTER
Zoha Laraib

AUTHORS
Zoha Laraib, Sherwood Richers, Alessandro Baroni, Kathleen Hamilton, and In-Saeng Suh.

ABSTRACT
The flavor transformation of neutrinos through quantum mechanics drives the explosion of a massive star and governs the composition of elements created in its ejected matter. These interacting systems permit neutrino dynamics due to many-body quantum correlations that develop oscillations on ‘fast’ timescales associated with entanglement entropy. The mean-field approximation is a compute-efficient way to model neutrino quantum mechanics on a classical computer, but it is still not clear how well it represents changes in flavor in a precise many-body treatment. We investigate this discrepancy through many-body simulations of neutrino quantum kinetics representing the supernova instabilities particularly on a quantum computer as quantum algorithms are adept at addressing large entanglement in this exponentially-growing Hilbert space, showcasing their power. Using quantum computing resources at ORNL, we simultaneously incorporate many-body entanglement in neutrinos, while addressing major domain-specific uncertainties like its impact on the fast flavor instability. In our study, we investigated the deviations we obtained in the data via simulating on various quantum hardware and compared our findings with classical results. These results could close a significant hole in the quantum kinetic theory of how neutrinos drive the explosion of massive stars and the generation of heavy elements in the universe.

Energy Burden Research

PRESENTER
Morgan Perry

AUTHORS
Morgan, Som, Daniel, Zhenglai

ABSTRACT
Poster presentation for OLCF - Summer Program

Computational Studies of Fluorine Impurities in Spent NMC Cathode Materials Toward Direct Recycling

PRESENTER
Zongtang Fang

AUTHORS
Zongtang Fang; Dianne Atienza Hay

ABSTRACT
Hydrothermal relithiation under oxidative conditions has been reported to be an efficient method to remove unavoidable fluorine impurities and rejuvenate cycle-aged cathode materials. The mechanism of fluorine impurity removal as well as relithiation is still unknown. In this work, hydrothermal relithiation in LiOH solution with H2O2 as an oxidative additive at 125 °C followed by calcination was able to fully recover the capacity of cycle-aged NMC532 cathode material from end-of-life commercial electric vehicle cells with a state-of-health of 75%. The adsorbed surface oxygen species from H2O2 act as catalysts to facilitate both relithiation and removing surface fluorine impurities on NMC532. Removal of transition metal fluoride in LiOH solution is a displacement reaction with an *-OH group replacing a *-F group. X-ray photoelectron spectroscopy and Raman spectroscopy combined with electronic structure calculations confirm the conversion of transition metal fluoride to lithium fluoride. The activation energy is reduced via the formation of a peroxide with the adsorbed oxygen to provide more reactive *-OH groups coupled with a redox process. The kinetics for relithiation and impurity removal with hydrothermal method can be optimized by modifying surface oxygen.

A linear scaling ab initio approach to the electronic transport in disordered alloys

PRESENTER
Markus Eisenbach

AUTHORS
Vishnu Raghuraman, Xiao Liang, Markus Eisenbach, Yang Wang, Mike Widom

ABSTRACT
Electrical conductivity or resistivity is a fundamental property of a material and is largely influenced by the chemical disorder in the material. In this project, we have implemented computational methods for calculating the electrical conductivity with the Kubo-Greenwood (KG) formula in the KUBO component in the MuST package, which is a separate CSSI funded project for building the framework of Green's function based ab initio methods. We have combined the KG formula with the locally self-consistent multiple scattering (LSMS) method, enabling us to calculate residual resistivity of large systems from the first-principles. As a linear scaling ab initio method, LSMS can perform efficient first-principles calculations of systems with a large number of atoms. We show applications of this KG+LSMS method to a set of disordered alloys. We also show an application of KUBO, combined with KKR-CPA method, in the calculation of thermoelectricity of random alloys.

Dependence of the Detected Core-CollapseSupernova Neutrinos on the Nuclear Equation of State in the DUNE Neutrino Detector

PRESENTER
Colter Richardson, Antigoni Georgiadou, and Evan Semenak

AUTHORS
Colter Richardson, Antigoni Georgiadou, Evan Semenak, Anthony Mezzacappa, and Bronson Messer

ABSTRACT
The next galactic core-collapse supernova will provide a unique view into the extreme environments produced in one of the most energetic events in the Universe. Utilizing new probes in the form of neutrinos and gravitational waves we will be able to peer past the stellar atmosphere into the hyper dense core of these dying stars. We present on the dependence of neutrino detections on the nuclear equation of state in this extreme environment. As the densities of this newly formed proto-neutron star reach around 1e14 grams per cubic centimeter, these events are laboratories for probing high energy nuclear physics. As the core of these massive stars collapse, neutrinos carry away 99% of the energy and are produced deep within the star. By studying the neutrinos detected from these events we can better constrain many aspects of physics, from dark matter candidates to neutron-neutron interactions. With the Deep Underground Neutrino Experiment set to begin taking data soon, we have to be prepared for the next CCSN. We study the neutrino signatures produced by the Chimera core-collapse supernova code's E-Series, where only the nuclear equation of state is varied across the 7 models. We use the SNOwGLoBES software to simulate the neutrino interactions in DUNE and investigate how the nuclear equation of state impacts the detected signal.

Dependence of the Reconstructed CCSN Gravitational Wave High-Frequency Feature on the Nuclear Equation of State, in Real Interferometric Data

PRESENTER
Robert Daniel Murphy

AUTHORS
R. Daniel Murphy, Alejandro Casallas-Lagos, Anthony Mezzacappa, Michele Zanolin, Ryan Landfield, Eric Lentz, Pedro Marronetti, Javier Antelis, and Claudia Moreno

ABSTRACT
We present an analysis of gravitational wave (GW) predictions from five two-dimensional Core Collapse Supernova (CCSN) simulations that varied only in the Equation of State (EOS) implemented. The GW signals from these simulations are used to produce spectrograms in the absence of noise, and the emergent high-frequency feature (HFF) is found to differ quantitatively between simulations. Below 1 kHz, the HFF is well approximated by a first-order polynomial in time. The resulting slope was found to vary between 10–50% from model to model. Further, using real interferometric noise we investigated the current capabilities of GW detectors to resolve these differences in HFF slope for a Galactic CCSN. We find that, for distances up to 1 kpc, current detectors have the ability to resolve HFF slopes differing by 4.4–15%. For further Galactic distances, current detectors are capable of distinguishing the upper and lower bounds of the HFF slope for groupings of our models that varied

in EOS. With the higher sensitivity of future GW detectors, and with improved analysis of the HFF, our ability to resolve properties of the HFF will improve for all Galactic distances. This study shows the potential of using the HFF of CCSN produced GWs to provide insight into the physical processes occurring deep within CCSN during collapse, and in particular its potential to further constrain the EOS through GW detection.

Thermal Conductivity Analysis of Phase Change Materials for Water Heaters Using Phonon Dispersion

PRESENTER
Gonzalo Vela

AUTHORS
Gonzalo Vela

ABSTRACT
Three phase change materials, substances that store or release thermal energy while melting of freezing, were studied to determine their thermal conductivity while being used in hot water heat pumps. Using a series of programs, the phonon dispersion graph for each substance was obtained from which predictions about the substance's thermal conductivity were made. For large water pumps a high thermal conductivity is ideal, allowing for a quick absorbance and transfer of heat.

Scalable Multi-Facility Workflows for Artificial Intelligence Applications for Earth Observation Datasets

PRESENTER
Takuya Kurihana

AUTHORS
Takuya Kurihana, Tyler J. Skluzacek, Rafael Ferreira da Silva, Valentine Anantharaj

ABSTRACT
Earth observation and earth system models are sources of vast, multi-modal datasets that are invaluable for advancing climate and environmental research. However, their scale and complexity pose challenges for processing and analysis. In this paper we discuss our experiences in developing a scientific research application using an automated multi-facility workflow that orchestrates data collection, preprocessing, artificial intelligence (AI) inferencing, and data movement across diverse computational resources, leveraging the Advanced Computing Ecosystem Testbed at the Oak Ridge Leadership Computing Facility (OLCF). We demonstrate that our AI application workflow can be seamlessly integrated and orchestrated across research facilities to extract new scientific insights from climate datasets using data intensive computational methods. The results indicate that the multi-facility workflow reduces processing time, enhances scalability, and maintains high efficiency across varying workloads. Our workflow processes 12,000 high-resolution satellite images in 44 seconds using 80 workers distributed across 10 nodes on the OLCF systems.

Leveraging GNU Parallel for Optimal Utilization of HPC Resources

PRESENTER
Ketan Maheshwari

AUTHORS
Ketan Maheshwari, Ahmad Maroof Karimi, Junqi Yin, Fred Suter, Seth Johnson and Rafael Ferreira da Silva

ABSTRACT
GNU Parallel is a versatile and powerful tool for process parallelization widely used in scientific computing. This poster demonstrates its effective application in high-performance computing (HPC) environments, particularly focusing on its scalability and efficiency in executing large-scale high-throughput high-performance computing (HT-HPC) workflows. We highlight GNU Parallel's performance across various HPC workloads, including GPU computing, container-based workloads, and node-local NVMe storage. Our results on two leading supercomputers, OLCF's Frontier and NERSC's Perlmutter, showcase GNU Parallel's rapid process dispatching ability and its capacity to maintain low overhead even at extreme scales. Beyond its direct application as a viable workflow manager, GNU Parallel can be employed in conjunction with other workflow systems as a ``last-mile" parallelizing driver and as a quick prototyping tool to design and extract parallel profiles from application executions. We present the argument that the potential for GNU Parallel to transform workflow management at extreme scales is substantial, paving the way for more efficient and effective scientific discoveries.

IRI Science Applications on the Advanced Computing Ecosystem Testbed

PRESENTER
David M. Rogers

AUTHORS
David Rogers, Gustav Jansen, Jens Glaser, Ketan Maheshwari, Ross Miller, Rick Archibald, Tyler Skluzacek, Ryan Prout, A. J. Ruckman, Paul Bryant, Michael Brim, Rafael Ferreira da Silva, Tom Beck and Sarp Oral

ABSTRACT
GNU Parallel is a versatile and powerful tool for process parallelization widely used in scientific computing. This poster demonstrates its effective application in high-performance computing (HPC) environments, particularly focusing on its scalability and efficiency in executing large-scale high-throughput high-performance computing (HT-HPC) workflows. We highlight GNU Parallel's performance across various HPC workloads, including GPU computing, container-based workloads, and node-local NVMe storage. Our results on two leading supercomputers, OLCF's Frontier and NERSC's Perlmutter, showcase GNU Parallel's rapid process dispatching ability and its capacity to maintain low overhead even at extreme scales. Beyond its direct application as a viable workflow manager, GNU Parallel can be employed in conjunction with other workflow systems as a ``last-mile" parallelizing driver and as a quick prototyping tool to design and extract parallel profiles from application executions. We present the argument that the potential for GNU Parallel to transform workflow management at extreme scales is substantial, paving the way for more efficient and effective scientific discoveries.

Power Outages

PRESENTER
Alyvia Bitterly

AUTHORS
Suzanne Parete-Koon, Alan Longcoy, and Johnny McCormick

ABSTRACT
I will be presenting my project about Power Outages and how I used zip code resolution to track and collect data from a derecho storm in Texas.

Accelerating Discovery with NAMD 3 Simulation

PRESENTER
David J. Hardy

AUTHORS
David J. Hardy, Eric J. Bohm, Haochuan Chen, Barry Isralewitz, Rafael Bernardi, Emad Tajkhorshid

ABSTRACT
“Molecular dynamics simulations serve as a bridge between structural data and biological function, offering an atomic-scale view of the mechanistic foundations of life. The latest release of NAMD, version 3, provides a significant leap in computational capability, pushing the envelope into the realm of million-to-billion atom simulations - domains critical for understanding complex cellular functions. In speeding up the software, a cornerstone of this advancement is a new GPU-resident operational mode, which substantially accelerates molecular simulation running as a single multi-threaded process on compatible AMD and NVIDIA devices. With NAMD 3, a modern GPU can achieve simulation speeds of one microsecond per day for systems of approximately 25,000 atoms. This remarkable efficiency is further magnified in GPU-dense single-node environments, where tightly integrated communication interconnects between GPUs enable scaling for larger systems, ranging from one million to twenty million atoms. Such capability enables simulation of biologically realistic systems at speeds of hundreds of nanoseconds per day. Key functionalities such as steered molecular dynamics (SMD), collective variables (Colvars), and alchemical free energy methods are now available in the new GPU-resident mode, addressing the evolving needs of the NAMD user community. NAMD supports launching ensemble simulations across multiple compute nodes, running each replica as a single GPU-resident process on single devices up to multi-GPU scaling on single nodes. NAMD still supports multi-node scaling of a single simulation running in a GPU-offload mode, and recent advances to NAMD’s Charm++ runtime system now provide improved scaling on OLCF Frontier and other Slingshot 11 installations.”

GPU-enabled Extreme Scale Turbulence Simulations: 32768^3 grid resolution and application to mixing and dispersion

PRESENTERS
Rohini Uma-Vaideswaran and Daniel L. Dotson

AUTHORS
Rohini Uma-Vaideswaran, Daniel L. Dotson, P.K. Yeung, K. Ravikumar, S. Nichols

ABSTRACT
Direct numerical simulations of 3D fluid turbulence with disorderly fluctuations spanning a wide range of length scales and time scales have long been a grand challenge problem in both leadership-class computational science. Our work since the beginning of the CAAR Program For Frontier has led to the successful development of a new GPU code with very favorable performance characteristics and capable of simulating turbulence at resolution up to 32768^3 resolution (35 trillion grid points), which is the largest known worldwide at this time. In this poster we first briefly summarize the essential elements of the prior work and the early science results that it has enabled, especially for studies of extreme events which are localized in time and space. Then we report on more recent algorithmic developments which are important for studies of turbulent mixing and dispersion. A particle-tracking algorithm based on cubic splines interpolation for the particle velocity scales very well, with spline coefficients generated using OpenMP offload, and a dynamic mapping between particles and parallel processes that faciliates large particle counts with little increase in overall cost. For the mixing of passive scalars the case of low molecular dffusivity (i.e. high Schmidt number) is well addressed by a separate algorithm that allows asynchronism between operations carried out by separate groups of parallel processes performing calculations at two different grid resolutions. All of the codes discussed show very good strong scaling and weak scaling, subject only to limitations from small message sizes at node counts approaching full machine capacity on Frontier. Finally we show some results on 3D visualization where large datasets also pose specific challenges.

Dust in the Wind: Simulation dust evolution in galactic outflows

PRESENTER
Helena Richie

AUTHORS
Helena Richie and Evan Schneider

ABSTRACT
“Astrophysical dust forms in the interstellar medium (ISM) of galaxies during the deaths of stars. Despite this, dust is frequently observed far from the ISM, in the haloes of galaxies. Galactic winds are one possible explanation for this since they are driven by supernova explosions in the ISM and could sweep clouds of cool, dusty ISM gas out of galaxies and into the halo. However, these winds’ hot, turbulent conditions can also destroy dust via sputtering. We aim to determine whether dust can survive in hot outflows by directly modeling dust transport and destruction in numerical simulations of galactic winds. We use a novel dust model in the Cholla hydrodynamics code to create these simulations. Cholla’s GPU-native nature allows us to simulate large physical volumes at fixed, field-leading resolution. Previously, we used this capability to create high-resolution, idealized simulations of individual dusty clouds accelerated by hot winds. By resolving the detailed mixing of dust in

cloud-wind interactions, we learned that cloud shielding can enable long-term dust survival in hot winds. Now, we build on this work by using this model to create the first high-resolution simulations of entire galaxies with dust, providing exciting new insight into the evolution of dust in galactic outflows."

Ensemble Simulations on Leadership Computing Systems

PRESENTER
Antigoni Georgiadou

AUTHORS
Antigoni Georiadou, Henry Camacho, Tanvir Sohail, Swarnava Ghosh, Arjun Valiya Parambathu, Dilipkumar N. Asthagiri, Dmytro Bykov, Tushar Athawale, and Thomas L. Beck

ABSTRACT
Scientific productivity can be enhanced through workflow management tools, relieving large High-Performance Computing (HPC) system users from the tedious tasks of scheduling and designing the complex computational execution of scientific applications. This paper presents a study on the usage of ensemble workflow tools to accelerate science using the Summit and Frontier supercomputing systems. The research aims to connect science domain simulations using Oak Ridge Leadership Computing Facility (OLCF) supercomputing platforms with ensemble workflow methods in order to accelerate HPC-enabled discovery and boost scientific impact. We present the coupling, porting and optimization of Radical-Cybertools on three applications: Chroma, NAMD and LAMMPS. The tools augment traditional HPC monolithic runs with a pilot scheduler. Lessons-learned are discussed for physics, biology and materials science applications. We discuss intrinsic limitations of coupling and porting ensemble workflow tools to applications that run on large HPC systems. The origins of technical challenges and their solutions developed during the implementation process are discussed. Data management strategies, OLCF's policies for ensembles, and natively supported workflow tools are also summarized.

Solving the Multiscale Mutliphysics Puzzles in ITER Edge

PRESENTER
CS Chang

AUTHOR
C.S. Chang, R.M. Churchill, J. Dominski, R. Hager, S. Ku, B. Sturdevant, G. Wilkie, S. Abbott, M. Adams, R. Bustamante, L. Chacon, Y. Chen, J. Choi, S. Ethier, A. Gainaru, A. Georgiadou, J. Gu, K. Huck, S. Klasky, P. Lin, S. Parker, N. Podhorszki, A. Scheinberg, M. Shephard, E. Suchyta, T. Williams, K. Wu

ABSTRACT
This poster is a reproduction of the Invited User Talk given by C.S. Chang. Experiments on present-day tokamaks have observed over the last 40 years that the edge plasma behaviors determine the core fusion performance in a critical way. However, many aspects of the edge plasma behaviors have been puzzling. They could not be easily modeled due to the nonlinear multiscale multi-physics nature of tokamak edge plasma that required first-principles simulation on extreme scale computers. Compounding the difficulties, the hot, dense, high-magnetic field, and large-size ITER could be in a different physics regime from the present-day tokamaks, which could make the experimental physics extrapolation unreliable.

Fortunately, world's largest capability-computers at OLCF enabled our XGC team to simulate tokamak multiscale multi-physics edge plasma behaviors via the INCITE and ECP programs. We have solved several edge physics puzzles and indeed found that ITER edge plasma is in a different physics regime with different edge plasma behaviors, which could help the operation of ITER and design of future magnetic fusion reactors such as FPP (Fusion Pilot Plant). XGC (X-point Gyrokinetic Code) is a modern particle-in-cell code designed for extreme scale GPU computers mostly written in C++. It uses first-principles-based kinetic equations and solves tokamak plasmas in realistic geometry that includes the magnetic separatrix, X-point and material wall boundary on unstructured triangular mesh. It has been an ECP, SciDAC, INCITE and Aurora-ESP code, utilizing Kokkos for performance portability. It shows an excellent scalability on Frontier, Perlmutter, and Aurora (up to the available node count of ~2,048). XGC's performance, especially on Frontier, is under deligent improvement through algorithm modification, data flow optimization, further utilization of GPU usage, and software adaptation.