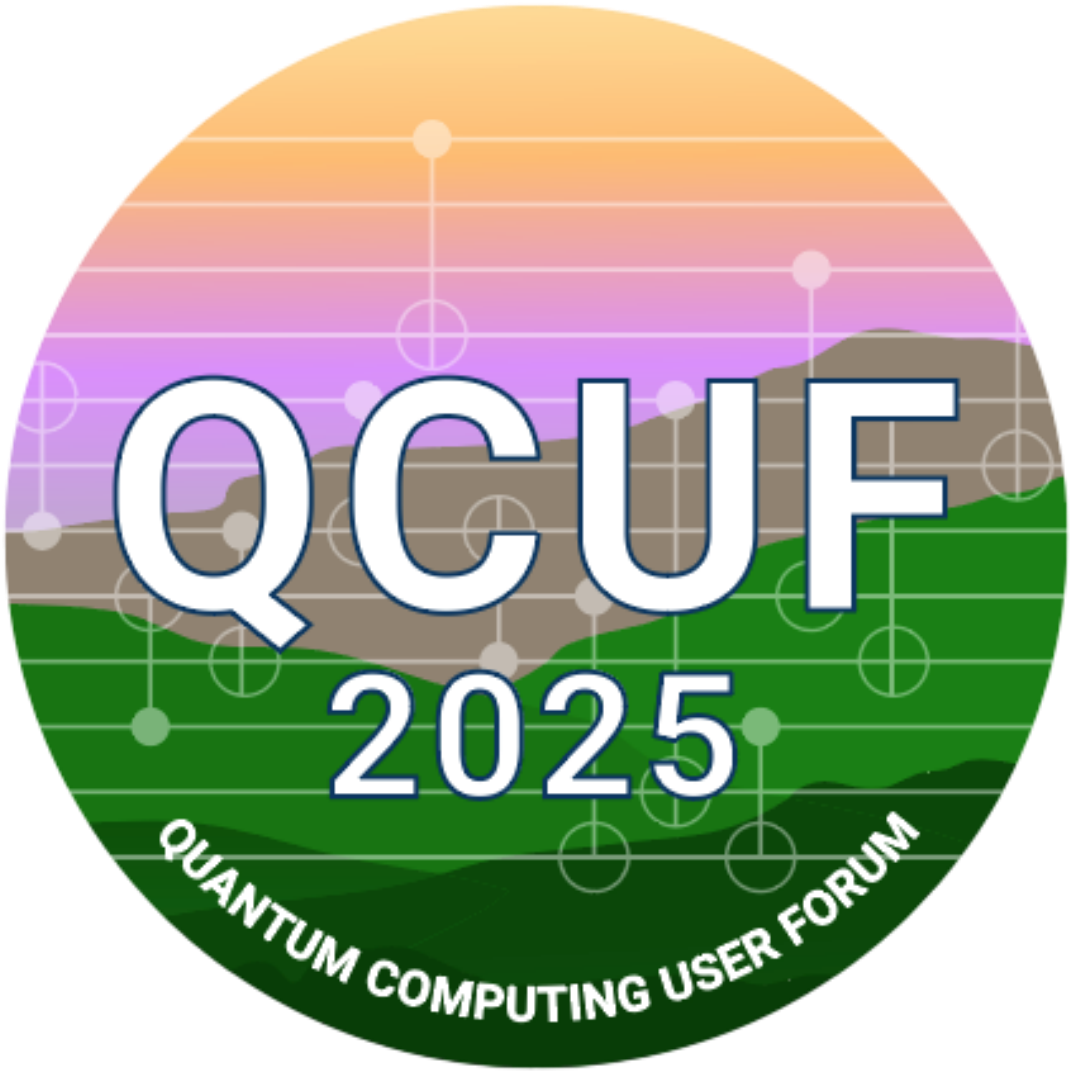


# Abstract Booklet

## Quantum Computing User Forum

## July 21-24th, 2025

*Hosted by the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory*

OAK RIDGE National Laboratory

# Table of Contents

## KEYNOTE TALKS

| Martin Roetteler | IonQ | 7/21, 1pm |
|---|---|---|

### TOWARD QUANTUM COMMERCIAL ADVANTAGE WITH ENTERPRISE-GRADE QUANTUM COMPUTERS

Quantum computers might achieve commercial advantage over classical computers for specific problems. The IonQ Applications team identifies applications across a broad spectrum of use cases, investigates the cross-over points with classical compute, and works closely with customers and partners to understand commercial use cases and the integration of quantum applications into larger workflows. In this presentation, we give an overview of areas where we believe that there are a few candidates for commercial advantage on near-term systems.
This includes:

- Chemistry
- Optimization
- Topological Data Analysis
- Quantum + AI

| Justin Bohnet | Quantinuum | 7/22, 1pm |
|---|---|---|

### HELIOS: A NEW DIMENSION FOR QUANTINUUM SYSTEMS

This year, Quantinuum will release the Helios generation of trapped-ion quantum computers, based on a new quantum charge-coupled device (QCCD) architecture. This release marks two major advances for Quantinuum systems, as Helios will be the first to use a 2D junction for qubit routing and the first to use 137Ba+ as qubits. These changes enable Helios to offer more qubits with lower physical error rates, all while operating faster than the current generation H2 systems. Combined with an all-new application programming language, these capabilities enable users to better explore the hybrid classical-quantum application space and solve more classically intractable problems. In this presentation, you will learn how Helios fits into our hardware roadmap, and how to prepare to get the most out of the new capabilities of the improved hardware and software stack.

| Sarah Sheldon | IBM | 7/23, 1pm |
|---|---|---|

**FROM ALGORITHMS TO APPLICATIONS: THE PATH TO QUANTUM ADVANTAGE**

Achieving quantum advantage will require building quantum hardware and software capable of running circuits at a scale beyond classical simulation, developing methods for extracting accurate measurements with high confidence, and identifying problems with structure that quantum algorithms can address. This talk will give an overview of IBM Quantum's progress towards each of these steps. The quantum systems available today and in the coming years allow us to explore applications of quantum computing at a scale that was previously inaccessible and to develop rigorous methods for evaluating heuristic quantum algorithms. As the technology matures, we invite a new period of scientific exploration that will rely on collaboration between quantum scientists and domain experts to realize useful quantum computing.

| Eric Mansfield | IQM | 7/24, 1pm |
|---|---|---|

**SUPERCHARGING HPC WITH IQM: A VISION FOR HYBRID COMPUTING**

IQM Quantum Computers is pioneering the integration of quantum computing (QC) with high-performance computing (HPC). One of the significant unanswered questions in the realm of HPC+QC workflows is how to pinpoint the workflows that yield the greatest interest and impact. We will present collaborative research done with our partners, highlighting the fields of chemistry simulation, optimization, and quantum machine learning. We will discuss why these workflows are ideal for HPC+QC and how the workflows were executed on IQM hardware. Finally, we present our vision for how HPC+QC will evolve in the NISQ era and what the first fault tolerant workflows could look like.

## CONTRIBUTED TALKS

### EVALUATING THE PERFORMANCE OF QUANTUM PROCESS UNITS AT LARGE WIDTH AND DEPTH

Presenter: Alejandro Montanez-Barrera
Affiliation: Forschungszentrum Julich

Currently, we are in a stage where quantum computers surpass the size that can be simulated exactly on classical computers, and noise is the central issue in extracting their full potential. Effective ways to characterize and measure their progress for practical applications are needed. In this work, we use the linear ramp quantum approximate optimization algorithm (LR-QAOA) protocol, a fixed quantum QAOA protocol, as an easy-to-implement scalable benchmarking methodology that assesses quantum process units (QPUs) at different numbers of qubits and 2-qubit gate depths. The benchmarking identifies the depth at which a fully mixed state is reached, and therefore, the results cannot be distinguished from those of a random sampler. We test this methodology using three graph topologies: 1D-chain, native layout, and fully connected graphs, on 19 different QPUs from 5 vendors: IBM, IQM, IonQ, Quantinuum, and Rigetti for problem sizes requiring from Nq=5 to Nq=156 qubits and LR-QAOA number of layers from p=3 to p=10,000. In the case of 1D-chain and native graphs, ibm_fez, the system with the largest number of qubits, performs best at p=15 for problems involving Nq=100 and Nq=156 qubits and 1,485 and 2,640 2-qubit gates, respectively. For the native graph problem, ibm_fez still retains some coherent information at p=200 involving 35,200 fractional 2-qubit gates. Our largest implementation is a 1D-chain problem with p=10,000 involving 990,000 2-qubit gates on ibm_fez. For fully connected graph problems, quantinuum_H2-1 shows the best performance, passing the test with Nq=56 qubits at p=3 involving 4,620 2-qubit gates with the largest 2-qubit gate implementation for a problem with Nq=50 qubits and p=10 involving 12,250 2-qubit gates but not passing the test.

### UTILITY SCALE IMPLEMENTATION AND BENCHMARKING OF DYNAMICALLY CORRECTED GATES AT THE COHERENCE LIMIT OF IBM DEVICES

Presenter: Hisham Amer
Affiliation: Virginia Tech

Utility-scale, gate-based, quantum computing relies primarily on maintaining high fidelity gates in a noisy environment, on many qubits and for an extended period of time. Here, we implement a control scheme at scale that simplifies this task, needing next to no calibration, by utilizing dynamically corrected gates (DCGs) where inherent noise suppression is baked into the control fields underlying the gates. To implement DCGs, we use our own Space Curve Quantum Control (SCQC) formalism which ties quantum evolution to the traversing of geometric space curves. Our SCQC gate design is powered by our publicly available Qurveros package which automates the generation of arbitrary single qubit unitaries through experimentally viable, noise-robust, DCGs that can be comfortably tailored to a device's practical constraints. To experimentally

demonstrate the utility of SCQC we implement, and extensively benchmark, a set of dynamically corrected {Hadamard, SX, X} gates on IBM's superconducting devices. We compare the general performance, including noise robustness, of IBM's native DRAG pulses to our DCGs on a set of 21 qubits achieving noise robust gates at error rates as low as 7e-5. Additionally, we demonstrate the conditions under which these gates can be made to operate at the coherence limit.

## EHANDS: QUANTUM PROTOCOL FOR POLYNOMIAL COMPUTATION ON REAL-VALUED ENCODED STATES

Presenter: Jan Balewski
Affiliation: NERSC/LBNL

The novel constructive EHands protocol defines a universal set of quantum operations for multivariable polynomial transformations on quantum processors by introducing four basic subcircuits—multiplication, addition, negation, and parity flip—and by using expectation-value encoding (EVEN) to represent real numbers in quantum states. These elementary arithmetic operations can be systematically composed to compute degree-d polynomials, P_d(x), on a QPU. The resulting quantum circuit structure closely mirrors the stepwise evaluation of polynomials on a classical calculator, providing an intuitive and efficient approach to polynomial computation on quantum hardware. By enabling direct, predictable polynomial and nonlinear data transformations on a QPU, our method reduces dependence on classical post-processing in hybrid quantum-classical algorithms, enabling advancements in many quantum algorithms. The EHands quantum circuits are compact enough to deliver meaningful and accurate results on today's noisy quantum processors. We present a detailed implementation of P_4(x) and report experimental results for polynomial approximations of common functions, obtained using IBM's Heron-class quantum processors and an ideal Qiskit simulator.

## MAGNETISM AND DYNAMICS USING QUANTUM HARDWARE

Presenter: Arnab Banerjee
Affiliation: Purdue University

Quantum magnetism is an emergent field which promises not only exotic emergent ground states, with phases hosting fermionic, bosonic and anyonic states, but also those which are useful for future quantum devices, qubits and sensors. Understanding these states in real materials is often hard. But it is likely that quantum computers performing quantum simulations on Hamiltonians would be able to shed light on not only the data analysis, but also provide a pathway towards new material designs. As we inch towards that goal, in this talk, I will touch upon three different problem classes that bring experiments on quantum magnets closer to quantum computers. The first one is on the simulation of the dynamic structure factor which is measured in neutron scattering. The second is the simulation of a quantum transport phenomena. The third one is the simulation of the Kardar Parisi Zhang hydrodynamics and its breakdown beyond the integrable limits which can help relate to experiments performed in spin chain compounds.

## QUANTUM SIMULATION OF COLOR IN HIGH-ENERGY QUANTUM CHROMODYNAMICS

Presenter: Herschel Chawdhry
Affiliation: Florida State University

Quantum Chromodynamics (QCD) describes the fundamental interactions of quarks and gluons, which are the constituent particles of protons and neutrons. To predict the behavior of these particles at high energies, such as in particle colliders, QCD calculations are vital but exceedingly challenging computationally. In this talk I will present first steps towards using quantum computers to model high-energy QCD processes. In particular, I present quantum computing algorithms to model the color part of high-energy QCD interactions, and show how they allow scattering amplitudes and quantum interferences to be manipulated in a very natural way on a quantum computer. Using ORNL QCUP resources, we demonstrate simplified examples of our algorithms on present-day quantum computing devices.

## APPLICATIONS OF QUANTUM COMPUTING IN POWER GRID/ENERGY SYSTEM - SOME EXAMPLES

Presenter: Suman Debnath, Phani Marthi
Affiliation: Oak Ridge National Laboratory

Solving large-scale power systems problems related to planning, and operations on classical computers have had several challenges due to their scalability and inherent nature (such as mixed-integer and nonlinear). Quantum computing is one of the emerging computing paradigms that is well suited to solve these challenges because of its potential quantum advantage over classical computers. To that end, in this talk, quantum algorithms targeted to solve different large-scale power systems problems will be briefly presented. A variational quantum algorithm (VQA)-based new hybrid quantum algorithm to solve the Unit Commitment (UC) problem in power systems will be presented in detail. UC is one of the power system mixed-integer nonlinear optimization problems where the objective is to minimize the total cost while optimally allocate the generating units to meet the hourly demand of the power loads. The effectiveness of the proposed quantum algorithm is presented through testing and evaluation of three different use-cases. The evaluation results of the proposed algorithm on IonQ quantum computers will also be presented. In addition to UC, use of variational quantum algorithms for EMT simulation and their challenges will be presented.

## UTILIZING QUANTUM COMPUTING FOR SOLVING MULTIDIMENSIONAL PARTIAL DIFFERENTIAL EQUATIONS

Presenter: Esam El-Araby
Affiliation: The University of Kansas (KU)

Multidimensional partial differential equations (PDEs) are fundamental to modeling various physical phenomena. Processing a large amount of information is a significant capability of quantum computing particularly for solving complex and/or multidimensional PDEs. There are currently many quantum techniques available for solving PDEs, which are mainly based on variational quantum circuits. However, the existing quantum PDE solvers, particularly those based on variational quantum eigensolver (VQE) techniques, suffer from several limitations. These include low accuracy, high execution times, and low scalability on quantum simulators as well as on noisy intermediate-scale quantum (NISQ) devices, especially for multidimensional PDEs. We, here, propose an efficient and scalable algorithm for solving multidimensional PDEs. We present two variants of our algorithm: the first leverages finite-difference method (FDM), classical-to-quantum (C2Q) encoding, and numerical instantiation, while the second employs FDM, C2Q, and column-by-column decomposition (CCD). Both variants are designed to enhance accuracy and scalability while reducing execution times. We have validated and evaluated our algorithm using the multidimensional Poisson equation and heat equation as case studies. Our results demonstrate higher accuracy, higher scalability, and faster execution times compared to VQE-based solvers on noise-free and noisy quantum simulators from IBM. Additionally, we validated our approach on hardware emulators and actual quantum hardware, employing noise mitigation techniques. We will also focus on extending these techniques to PDEs relevant to computational fluid dynamics and financial modeling, further bridging the gap between theoretical quantum algorithms and practical applications.

## QUANTUM LINEAR SOLVERS FOR COMPUTATIONAL FLUID DYNAMICS

Presenter: Kalyana Chakravarthi Gottiparthi
Affiliation: Oak Ridge National Laboratory

Fluid flow is ubiquitous in nature and often complex due to interplay of turbulence, compressibility, and chemical reactions. Computational simulations have become indispensable and provide valuable insights into the physics of fluids by resolving a wide range of length and time scales. Although advances in computing have enabled high-fidelity numerical simulations, they are still prohibitively expensive for many applications. Recent progress in quantum computing (QC) provides an opportunity to address challenges in computational fluid dynamics (CFD) and expand the current state-of-the-art in accuracy and speed of computations. However, quantum computing algorithms and hardware capable of handling CFD applications are still in early stages of development and require improvements to be useful in problems of interest. In order to utilize and improve on the advancements made in quantum computing for CFD, we implemented a fluid flow solver based on quantum algorithms for linear system of equations. We addressed canonical ideal flows using Harrow-Hassidim-Llyod (HHL) algorithm and assessed the errors along with some error mitigation strategies. The scalability of the implementation on different simulators, emulators, and quantum hardware has been analyzed. In this talk, we will discuss these results from our current efforts along with improvements to the quantum algorithm which enabled to scale the qubits to problems of practical interest. We will also discuss successful outcomes from our Quantum Computing User Program (QCUP) allocation such as convergence

of High Performance Computing (HPC) and QC, workforce development via student competition, and workflow for hybrid QC and HPC applications.

## IMPROVED IMPLEMENTATION OF QUANTUM LINEAR SYSTEMS ALGORITHMS ON A TRAPPED-ION QUANTUM PROCESSOR

Presenter: Adrian Harkness
Affiliation: Lehigh University

Iterative Refinement is a classical post-processing scheme that can exponentially improve the accuracy to which a linear system of equations can be solved using low-precision arithmetic. In the context of quantum linear systems algorithms, such as the HHL algorithm proposed by Harrow, Hassidim, and Loyd, Iterative Refinement can greatly reduce the quantum resources required to calculate an accurate solution in terms of tomography cost, circuit volume, and fault-tolerant overhead. In this talk, we discuss our ongoing work to calculate highly accurate solutions to linear systems of equations of up to eight variables by running HHL with Iterative Refinement on a NISQ trapped-ion quantum computer. We also present our open-source implementation, emphasizing that our circuit is not tailored to specific problem instances as most available implementations are.

## MANY-QUBIT IS ALL WE NEED FOR OPTIMIZATION CHALLENGES

Presenter: Seongmin Kim
Affiliation: ORNL

Hybrid quantum-classical approaches show promise for optimization problems but face challenges in accuracy, scalability, and flexibility in the current quantum era. To address these issues, we develop a highly efficient ansatz that reduces the circuit complexity by relying solely on superposition for a variational quantum-classical algorithm. This ansatz allows interactions between variables to be effectively captured through parameter optimization like classical neural networks. By optimizing hyperparameters, this algorithm achieves high accuracy. Furthermore, we have observed speedup when leveraging a quantum device over simulator for this algorithm. While the simulator exhibits exponential time complexity for problem size n (when circuits are not partitioned), the quantum device maintains constant time complexity, emphasizing its efficiency for handling wide quantum circuits. Our algorithm utilizes distributed execution on high-performance computers to explore large solution spaces, consistently finding ground truth solutions for optimization problems and outperforming classical algorithms. It demonstrates exceptional performance across various applications, including classical binary optimization, quantum chemistry, and real-world problems. Notably, it accelerates material optimization tasks, such as metamaterial design, by over 50x compared to state-of-the-art optimization algorithms while achieving superior results by learning multi-variable interactions. Beyond binary optimization, our approach enables multi-state optimization by mapping N-ary variables onto the Bloch sphere. These capabilities position it as a highly promising and versatile solver for real-world optimization challenges, highlighting the practical utility of the quantum-classical approach.

## VARIATIONAL QUANTUM SIMULATIONS OF A TWO-DIMENSIONAL FRUSTRATED TRANSVERSE-FIELD ISING MODEL ON A TRAPPED-ION QUANTUM COMPUTER

Presenter: Ammar Ahmad Kirmani
Affiliation: LANL

Quantum computers are an ideal platform to study the ground state properties of strongly correlated systems due to the limitation of classical computing techniques particularly for systems exhibiting quantum phase transitions. While the error rates of Noisy Intermediate-Scale Quantum (NISQ) computers are still high, simulating strongly correlated systems on such devices and extracting information of possible phases may be within reach. The frustrated transverse-field Ising model (TFIM) is such a system with multiple ordered magnetic phases. In this work, we simulate a two-dimensional frustrated TFIM with next-nearest-neighbor spin-exchange interactions at zero temperature. The competition between the nearest-neighbor and next-nearest-neighbor couplings along with the presence of quantum fluctuations makes the ground-state phase profile rich. We use the Variational Quantum Eigensolver (VQE) to compute the phases on a square lattice with periodic boundary conditions for a system of 16 sites (qubits). The VQE circuits are executed on the Quantinuum H1-1 trapped-ion quantum computer without using any error mitigation techniques. Our experiments show near perfect recovery of the magnetic phases of the frustrated model through ground-state energy, the energy derivative, and the spin correlation functions. Thus, we show that the trapped-ion quantum processor is able to achieve reliable simulations of a strongly correlated system within the limitations of the VQE approach.

## STQS: A UNIFIED SYSTEM ARCHITECTURE FOR SPATIAL TEMPORAL QUANTUM SENSING

Presenter: Chenxu Liu
Affiliation: PNNL

Quantum sensing (QS) exploits quantum phenomena to measure physical observables with exceptional precision and sensitivity. As QS applications become more complex, systematically exploring distributed architectures is crucial. We present STQS, a unified spatiotemporal framework that integrates sensing, memory, communication, and computation. Using a gate-based model, we analyze noise at each stage of sensing through simulation and introduce a novel distance-based metric for comparing reference and sensing states. This measure assigns confidence levels and could lead to advanced quantum signal processing, including quantum machine learning. We demonstrate STQSâ€™s versatility with quantum object detection and qubit-based dark matter detection, validated on IBMâ€™s Marrakesh and IonQâ€™s Forte devices.

## PAULI CHECK EXTRAPOLATION FOR QUANTUM ERROR MITIGATION

Presenter: Ji Liu
Affiliation: Argonne National Laboratory

Pauli Check Sandwiching (PCS) is an error mitigation scheme that uses pairs of parity checks to detect errors in the payload circuit. While increasing the number of check pairs improves error detection, it also introduces additional noise to the circuit and exponentially increases the required sampling size. To address these limitations, we propose a novel error mitigation scheme, Pauli Check Extrapolation (PCE), which integrates PCS with an extrapolation technique similar to Zero-Noise Extrapolation (ZNE). However, instead of extrapolating to the 'zero-noise' limit, as is done in ZNE, PCE extrapolates to the 'maximum check' limit—the number of check pairs theoretically required to achieve unit fidelity. In this study, we focus on applying a linear model for extrapolation and also derive a more general exponential ansatz based on the Markovian error model. We demonstrate the effectiveness of PCE by using it to mitigate errors in the shadow estimation protocol, particularly for states prepared by the variational quantum eigensolver (VQE). Our results show that this method can achieve higher fidelities than the state-of-the-art Robust Shadow (RS) estimation scheme, while significantly reducing the number of required samples by eliminating the need for a calibration procedure. We validate these findings on both fully-connected topologies and simulated IBM hardware backends.

## OPTIMIZING QUANTUM LINEAR SOLVERS: BENCHMARKING THE HHL ALGORITHM WITH AN ENHANCED QPE FRAMEWORK

Presenter: Chao Lu
Affiliation: Oak Ridge National Laboratory

Harrow-Hassidim-Llyod (HHL) algorithm is a prominent quantum algorithm that offers exponential speedup over its classical counterparts by solving linear systems of equations with logarithmic complexity. However, synthesizing and executing HHL circuits demand significant computational resources from both classical and quantum systems. In this work, we benchmark the HHL algorithm using the optimized Quantum Phase Estimation (QPE) generation algorithm, LuGo, to enhance its scalability and efficiency. We leveraged Frontier and Perlmutter supercomputers to evaluate the scalability of generating HHL circuits and to measure the simulation time of the generated circuits. Additionally, we provide a comprehensive analysis of quantum hardware performance, including studies on qubit connectivity, fidelity comparisons, and average execution times. Our results offer preliminary insights into potential applications of the HHL algorithm enhanced by the LuGo framework and the performance on quantum hardware. Reference: Lu, C., Gopalakrishnan Meena, M., & Gottiparthi, K. C. (2025). LuGo: an Enhanced Quantum Phase Estimation Implementation. arXiv preprint arXiv:2503.15439 (https://doi.org/10.48550/arXiv.2503.15439).

## SPARSE NON-MARKOVIAN NOISE MODELING OF TRANSMON-BASED MULTI-QUBIT OPERATIONS

Presenter: Gregory Quiroz
Affiliation: Johns Hopkins University Applied Physics Laboratory

The influence of noise on quantum dynamics is one of the main factors preventing current quantum processors from performing accurate quantum computations. Sufficient noise characterization and modeling can provide key insights into the effect of noise on quantum algorithms and inform the design of targeted error protection protocols. However, constructing effective noise models that are sparse in model parameters, yet predictive can be challenging. In this work, we present an approach for effective noise modeling of multi-qubit operations on transmon-based devices. Through a comprehensive characterization of seven devices offered by the IBM Quantum Platform, we show that the model can capture and predict a wide range of single- and two-qubit behaviors, including non-Markovian effects resulting from spatio-temporally correlated noise sources. The model's predictive power is further highlighted through multi-qubit dynamical decoupling demonstrations and an implementation of the variational quantum eigensolver. As a training proxy for the hardware, we show that the model can predict expectation values within a relative error of 0.5%; this is a 7x improvement over default hardware noise models. Through these demonstrations, we highlight key error sources in superconducting qubits and illustrate the utility of reduced noise models for predicting hardware dynamics.

## COLLECTIVE NEUTRINO OSCILLATIONS IN THREE FLAVORS ON QUBIT AND QUTRIT PROCESSORS

Presenter: Ermal Rrapaj
Affiliation: NERSC, Lawrence Berkeley Laboratory

Collective neutrino flavor oscillations are of primary importance in understanding the dynamic evolution of core-collapse supernovae and subsequent terrestrial detection, but also among the most challenging aspects of numerical simulations. This situation is complicated by the quantum many-body nature of the problem due to neutrino-neutrino interactions which demands a quantum treatment. An additional complication is the presence of three flavors, which often is approximated by the electron flavor and a heavy lepton flavor. In this work, we provide both qubit and qutrit encodings for all three flavors, and develop optimized quantum circuits for the time evolution and analyze the Trotter error. We conclude our study with a hardware experiment of a system of two neutrinos with superconducting hardware: the IBM Torino device for qubits and Advanced Quantum Testbed device at the Lawrence Berkeley National Laboratory for qutrits. We find that error mitigation greatly helps in obtaining a signal consistent with simulations. While hardware results are comparable at this stage, we expect the qutrit setup to be more convenient for large-scale simulations since it does not suffer from probability leakage into nonphsycial qubit space, unlike the qubit setup.

## UNLOCKING QUANTUM POTENTIAL: FROM MOBILE ACCELERATORS TO SCALABLE PARALLEL CLUSTERS

Presenter: Andreas Sawadsky
Affiliation: Quantum Brilliance

As quantum computing evolves, the need for accessible, scalable, and deployable quantum systems becomes increasingly urgent. In this talk, I will share Quantum Brilliance's mission to revolutionize the quantum computing landscape by developing mobile quantum accelerators that can be deployed anywhere—from the edge to the cloud. We are building compact, mobile systems capable of running quantum algorithms in real-world environments, with the flexibility to operate alongside classical systems for hybrid parallelized applications. I'll highlight our vision by showcasing key milestones towards a miniatured quantum computer operation at room-temperature. I will also explore the parallelization potential of small quantum accelerators, detailing how multiple units can be deployed and coordinated for specific applications—offering enhanced computational power and flexibility for both scientific and industrial problems. Finally, I'll discuss our exciting partnership with ORNL where we explore integration, hybridization and parallelization of quantum computing clusters within classical HPC.

## BRIDGING PARADIGMS: DESIGNING FOR HPC-QUANTUM CONVERGENCE

Presenter: Amir Shehata
Affiliation: Oak Ridge National Laboratory

This talks presents a comprehensive software stack architecture for integrating quantum computing (QC) capabilities with High-Performance Computing (HPC) environments. While quantum computers show promise as specialized accelerators for scientific computing, their effective integration with classical HPC systems presents significant technical challenges. We propose a hardware-agnostic software framework that supports both current noisy intermediate-scale quantum devices and future fault-tolerant quantum computers, while maintaining compatibility with existing HPC workflows. The architecture includes a quantum gateway interface, standardized APIs for resource management, and robust scheduling mechanisms to handle both simultaneous and interleaved quantum-classical workloads. Key innovations include: (1) a unified resource management system that efficiently coordinates quantum and classical resources, (2) a flexible quantum programming interface that abstracts hardware-specific details, (3) A Quantum Platform Manager API that simplifies the integration of various quantum hardware systems, and (4) a comprehensive tool chain for quantum circuit optimization and execution. This talk provides a foundational blueprint for integrating QC capabilities into existing HPC infrastructures, addressing critical challenges in resource management, job scheduling, and efficient data movement between classical and quantum resources.

## DISTRIBUTED QAOA ON THE SUPERCONDUCTING AND TRAPPED-ION QUANTUM COMPUTERS

Presenter: In-Saeng Suh
Affiliation: ORNL

Quantum Approximate Optimization Algorithm (QAOA) is a promising heuristic for solving combinatorial optimization problems on near-term quantum devices. However, the scalability of QAOA is constrained by hardware limitations such as qubit connectivity, coherence time, and gate fidelity. In this work, we present a distributed implementation of QAOA that decomposes large Quadratic Unconstrained Binary Optimization (QUBO) problems into subproblems, which are solved independently on multiple quantum processors and integrated using classical post-processing techniques. Our approach leverages the superconducting and the trapped-ion quantum processors, demonstrating the feasibility of distributed QAOA across different hardware architectures. We compare performance metrics, including solution quality, circuit depth, and execution time, across both platforms and analyze the impact of qubit noise and connectivity constraints on optimization outcomes. Our results highlight the advantages and challenges of distributed QAOA and provide insights into hybrid quantum-classical strategies for large-scale optimization problems. Collaborators: ORNL: Seongmin Kim, Antonio Coello-Perez, Alessandro Baroni, Amir Shehata, Ryan Landfield, Michael Sandoval, Josh Cunningham, Tom Beck IonQ: Erica Stump, Heidi Nelson-Quillin IBM: Vincent Pascuzzi

## PROTEIN STRUCTURE PREDICTION IN DRUG DISCOVERY ON UTILITY-LEVEL QUANTUM PROCESSORS

Presenter: Yuqi Zhang
Affiliation: Kent State University

Quantum computing offers groundbreaking possibilities for simulating and optimizing complex systems. This study demonstrates an application of the Variational Quantum Eigensolver (VQE) to highly complex biomolecular problems. We propose an efficient quantum computing workflow that integrates a vector-based tetrahedral lattice modeling approach, Hamiltonian design based on the physicochemical properties of proteins, and a two-stage noise-adaptive workflow combining estimation and sampling, to predict the three-dimensional structure of protein pocket regions in ligand docking processes. By optimizing a quantum-classical hybrid framework, this method achieves efficient variational quantum optimization and highly accurate protein structure prediction and has been successfully validated on IBM's utility-level quantum processors. Our method is the first to solve practical protein structure prediction problems on utility-level quantum computing devices, demonstrating high generalizability and scalability, and fully validating the potential of quantum computing methods to address real-world problems. Experimental results show that, compared to state-of-the-art deep learning algorithms, our quantum method excels in key metrics such as ligand-receptor affinity and geometric stability in docking tasks. The upper and lower bounds of the RMSD(root-mean-square deviation) distribution for docking structures predicted by the quantum method were reduced by approximately 20%, and binding affinity scores improved by about 10%. Furthermore, this study demonstrates the feasibility of using quantum devices with limited resources to predict larger

protein structures relevant to drug discovery. We successfully predicted the full structure of the Abeta-42 protein on a utility-level quantu

## BYPASSING OPTIMIZATION: GENERATOR-COORDINATE-INSPIRED METHODS FOR QUANTUM CHEMISTRY

Presenter: Muqing Zheng
Affiliation: Pacific Northwest National Laboratory

Inspired by the generator coordinate method, we employ Unitary Coupled Cluster (UCC) excitation generators to construct non-orthogonal, overcomplete many-body bases, projecting the system Hamiltonian into an effective Hamiltonian, which bypasses issues such as barren plateaus that heuristic numerical minimizers often encountered in standard variational quantum eigensolver (VQE). Our method diverges from conventional quantum subspace expansion methods by introducing an adaptive scheme that robustly constructs the many-body basis sets from a pool of the UCC excitation generators. This approach supports the development of a hierarchical ADAPT quantum-classical strategy, enabling a balanced interplay between subspace expansion and ansatz optimization to address complex, strongly correlated quantum chemical systems cost-effectively, setting the stage for more advanced quantum simulations in chemistry. We also conducted experiments on IBM quantum computers, showing satisfying results with error mitigation techniques.

## LUNCHTIME TALKS

### THE QUANTUM COMPUTING USER PROGRAM

Presenter: Claire Marvinney and Josh Cunningham
Affiliation: ORNL
The Quantum Computing User Program (QCUP) advances the testing and evaluation of quantum computing technology by engaging the quantum computing community in competitive access to commercial quantum computing systems. The program supports a broad set of users and developers of quantum computing applications as well as multiple quantum technology vendors to further the impact of quantum computing on the US Department of Energy missions of scientific discovery and economic innovation. QCUP has continued to have a high impact on research for the quantum computing community by enabling access to the latest commercial quantum computing resources.

### QUANTUM COMPUTING FOR THE BENEFIT OF HUMANITY

Presenter: Catherine Lefebvre
Affiliation: Open Quantum Institute (OQI)
The Open Quantum Institute (OQI) is a multilateral governance initiative that promotes global and inclusive access to quantum computing and the development of applications for the benefit of humanity. As a novel science diplomacy instrument, it brings together research, diplomacy, private sector and philanthropy stakeholders. Initiated by the Geneva Science Diplomacy Anticipator (GESDA), OQI is now hosted by CERN for its pilot phase (2024-2026). OQI has four main activities: (1) accelerating applications for humanity; (2) promoting global access; (3) advancing capacity building; (4) activating multilateral governance. While quantum computing is still in its early stages of development and computational resources remain limited, there is an opportunity today to join a global effort to explore potential applications of the technology that will positively impact our society and our planet. Through the methodology and support of OQI, quantum experts and subject matter experts from around the world have been collaborating with UN agencies and large NGOs to explore the potential of quantum computing to address global challenges. OQI's use case portfolio include for instance use cases on leak detection in urban water distribution network, discovery of new antimicrobial compounds, editing plants genomics to increase crop yield, optimizing turbine layout in windfarms, etc.

### SUPPORTING EXASCALE SCIENCE AT ORNL: A BRIEF OVERVIEW OF THE OLCF

Presenter: Veronica Vergara
Affiliation: ORNL
The Oak Ridge Leadership Computing Facility (OLCF), a US Department of Energy Office of Science User Facility, accelerates scientific discovery by providing researchers worldwide with expert

support and access to some of the world's most powerful supercomputers. Simulations on OLCF systems—most notably the Summit and exascale Frontier supercomputers—have advanced reactor and aircraft safety, accelerated drug discovery, and deepened understanding of complex physical phenomena. The OLCF also leads quantum computing efforts and fosters collaboration through robust user programs, including QCUP, and interactive events. By equipping researchers with leadership-class resources and expertise, the OLCF drives scientific innovation on a global scale.

## VIRTUAL QPUS: BRIDGING THE GAP BETWEEN SIMULATION AND HARDWARE PERFORMANCE

Presenter: Andreas Sawadsky
Affiliation: Quantum Brilliance

In this talk, we'll explore Quantum Brilliance's Virtual QPU (vQPU), which offers a cost-effective and efficient way to simulate quantum hardware performance. By accurately mimicking Quantum Brilliance's hardware Quantum Development Kit (QDK), including noise models, vQPUs provide a more realistc experience to actual quantum hardware, making them an invaluable tool for testing, parallelization, and application development. We'll discuss how vQPUs allow users to explore parallelization strategies and integrate quantum solutions without the high cost and logistical challenges of physical hardware. With realistic noise models, vQPUs replicate the true behaviour of quantum devices, enabling more accurate testing and application discovery. I'll showcase examples where vQPUs have been used to demonstrate the power and flexibility of virtual quantum processing in action.

# POSTERS

## GRAPH MATCHING TROTTERIZATION FOR EFFICIENT CIRCUIT-BASED CONTINUOUS-TIME QUANTUM WALK SIMULATION

Presenter: Mostafa Atallah
Affiliation: The University of Tennessee, Knoxville

Continuous-time quantum walks (CTQWs) are a universal method of computation with a wide range of applications. There exist efficient circuit-model implementations of CTQWs for select graphs; however, little is known about their efficient implementations for arbitrary graphs. In this work, we present a method to approximately implement a CTQW on any graph in the circuit model. Our approach decomposes the original CTQW graph into edge matchings and utilizes these matchings as a basis for Trotterization. This approach extends to general Hamiltonian simulation applications, providing a scalable framework for implementing graph-based quantum algorithms on gate-model quantum computers.

## LINEAR RESPONSE VIA QUANTUM PHASE ESTIMATIONN

Presenter: Alessandro Baroni
Affiliation: ORNL

Quantum phase estimation can be used to obtain linear response functions in the frequency domain, a quantity challenging to calculate for various many-body quantum systems. In this poster, we present the application of this quantum algorithm to extract the linear response for a nuclear two-body system. We present preliminary results for Hubbard model-like systems.

## ENTANGLEMENT BENCHMARKING IN QUANTUM SIMULATIONS OF SPIN SYSTEMS

Presenter: Anshumitra Baul
Affiliation: Oak Ridge National Laboratory

Classical computers cannot handle the complexity of highly entangled states, limiting our understanding of quantum phenomena in condensed matter systems. Matrix Product States (MPS) offer a compact and efficient representation of quantum many-body systems—particularly in one dimension—by capturing key entanglement features inherent in quantum materials. Leveraging MPS to prepare highly entangled ground states on a quantum computer enables accurate modeling of these systems and facilitates the simulation of their evolution under varying temperatures and external fields. Furthermore, analyzing entanglement witnesses within these states provides insight into quantum phase transitions, as demonstrated by Laurell et al., PhysRevLett.127.037201.5. Ultimately, the aim is to harness quantum computers to simulate ground states, explore system dynamics, and investigate properties of quantum materials that are otherwise intractable with classical methods.

## THE QUANTUM-ACCELERATED INTERNET TESTBED

Presenter: Ashlyn Burch
Affiliation: Oak Ridge National Laboratory

As quantum computing continues to advance, distributed quantum computing is becoming increasingly necessary to leverage the best resources across different platforms. Central to this integration is the development of scalable quantum networks. The Quantum-Accelerated Internet Testbed (QuAInT) at Oak Ridge National Laboratory (ORNL) is a multi-institutional collaborative initiative designed to lay the groundwork necessary for a national quantum internet connection by enhancing intercity ground-based fiber-optic and free-space quantum networks as well as improving technology for satellite-based communication. In this poster, we explore some of the testbed's key components and progress including that of a solid-state quantum memory, quantum repeaters, and broadband frequency conversion, among others. We will also discuss our ideas to integrate and connect other emerging quantum technologies available at ORNL within the network.

## BENCHMARKING QUANTUM SIMULATORS FOR DQAOA AND TFIM USING QFW ON FRONTIER

Presenter: Srikar Chundury
Affiliation: NCSU (and GRO intern at ORNL)

Simulating quantum circuits is essential for validating algorithms like the Distributed Quantum Approximate Optimization Algorithm (DQAOA) and Transverse-Field Ising Model (TFIM) evolution. However, no single simulator consistently performs best– efficiency depends on circuit structure, entanglement, and depth.

In this poster, we integrate QTensor, a tensor network-based simulator, into the Quantum Framework (QFw)– a modular platform that supports multiple simulation backends, including Qiskit-Aer, NWQ-Sim, and TN-QVM, via a unified interface. QFw decouples quantum applications from simulator implementations, enabling seamless backend switching, performance benchmarking, and scalable execution on HPC systems with minimal code changes. We benchmark DQAOA and TFIM circuits across supported simulators, showing how performance varies significantly with problem type. All simulations are deployed on the Frontier supercomputer using QFw's MPI-based orchestration for distributed, multi-node execution. These results underscore the need for simulator-agnostic infrastructure to enable systematic evaluation and high-performance scaling of quantum workloads. QFw provides a practical and extensible path toward reproducible quantum algorithm development across diverse application domains.

## FAST-FORWARDED CHEMICAL DYNAMICS ON QUANTUM COMPUTERS

Presenter: Joshua M Courtney
Affiliation: University of Georgia

Nonadiabatic dynamics are brought through dynamical coupling between atomic and molecular degrees of freedom. With modern classical methods like multiconfiguration time-dependent Hartree, dynamics can be performed with long computational times and large memory requirements. Decomposition of a Hamiltonian describing molecular dynamics leads to polynomial complexity in memory and time on quantum computers. With a real-space representation of wavepacket dynamics, existing simulations of Marcus model dynamics are confirmed via a Trotterized Hamiltonian on a simulated fault-tolerant quantum computer. Post-processing the measurement of an ancillary qubit storing population fraction data enables the calculation of rate coefficients. With full knowledge of the system's Hamiltonian, a violation of the computational time-energy uncertainty relation allows for compression of Hamiltonian terms. Using a quantum-assisted quantum compiling prescription to minimize cost, terms of the Trotterized Hamiltonian are mapped to diagonal variational ansatze with the aim to still produce rate coefficients. Free and single-potential wavepacket propagations are shown to be fast-forwarded via a single ansatz circuit that is constant in number of Trotter steps. Variational ansatze coupled to ancillary and additional computational registers allow for faster chemical dynamics simulations
with structured sparse Hamiltonians. Resulting circuit depth scales logarithmically with computational basis size, Hamiltonian sparseness, number of potentials, and dimension, and linearly with number of Trotter steps.

## UTILIZING QUANTUM COMPUTING FOR SOLVING MULTIDIMENSIONAL PARTIAL DIFFERENTIAL EQUATIONS

Presenter: Esam El-Araby
Affiliation: The University of Kansas (KU)

Quantum computing has the potential to significantly transform computational problem-solving. It leverages the quantum mechanical principles of superposition and entanglement to process a large amount of information simultaneously. This capability is particularly important for the numerical solution of complex and/or multidimensional partial differential equations (PDEs). There are currently many quantum techniques available for solving PDEs, which are mainly based on variational quantum circuits. However, the existing quantum PDE solvers, particularly those based on variational quantum eigensolver (VQE) techniques, suffer from several limitations, including low accuracy, high execution times, and low scalability on quantum simulators as well as on noisy intermediate-scale quantum (NISQ) devices.

In this work, we propose an efficient and scalable algorithm for solving multidimensional PDEs. We present two variants of our algorithm: the first leverages finite-difference method (FDM), classical-to-quantum (C2Q) encoding, and numerical instantiation, while the second employs FDM, C2Q, and column-by-column decomposition (CCD). Both variants are designed to enhance accuracy and scalability while reducing execution times. We have validated and evaluated our

proposed concepts using several case studies including multidimensional Poisson equation, multidimensional heat equation, Black Scholes equation, and Navier-Stokes equations for computational fluid dynamics (CFD). Our results demonstrate higher accuracy, higher scalability, and faster execution times compared to VQE-based solvers on noise-free and noisy quantum simulators from IBM. We also validated our approach on hardware emulators and actual quantum hardware, achieving promising results. This work establishes a practical and effective approach for solving PDEs using quantum computing for engineering and scientific applications.

## SIMULATING NOISY QUANTUM SYSTEMS AT SCALE

Presenter: Amit Jamadagni Gangapuram
Affiliation: Oak Ridge National Laboratory

Noise remains inherent to current day quantum architectures of various kinds, thereby posing a challenge, not only to benchmark their performance but also to realize quantum algorithms and applications at scale. Tensor networks have emerged as efficient tools that effectively characterize quantum correlations in certain regimes. However, a collaborative approach to extend these tools to characterize and benchmark noisy quantum hardware remains largely underdeveloped. To this extent, in this talk, I will introduce a tensor network framework capable of efficiently characterizing noisy quantum systems. Further, I will present algorithms that scale on classical architectures thereby providing a means not only to tightly bound the limits of classical simulation but also gain insights into the potential regime of quantum advantage. As examples, I will present numerical simulations involving the evolution of a quantum system in a purely dissipative setting as well as a noisy evolution i.e., involving both unitary and dissipative noisy channels. Towards the end, I will discuss potential extensions to the above algorithms while also briefly presenting the avenues that are likely to benefit using the introduced framework.

## QUANTUM HAMILTONIAN PHYSICS ON AMD GPUS

Presenter: Antigoni Georgiadou
Affiliation: ORNL

Forefront applications such as the many-body theory scientific domain are big sparse systems applications. The vision of retrieving information from the Hamiltonian matrix presents a challenge when solving for eigenvalues on a very large sparse set of equations to develop new understanding of the full range of complex quantum systems. Overall, solvers need to be robust and distributed tools to retrieve the eigenvalues while the size of the matrix does not matter as much as the sparsity of the system. In our applications we expect that iterative solvers are not applicable; nor dense solvers. Applications can produce matrix sizes that approximate order of magnitude close to millions and billions for the representation of basis states. Our specialized libraries to be used on Frontier and its hardware are PETSc/SLEPc. The Scalable Library for Eigenvalue Problem Computations (SLEPc) is based on the PETSc data structures with PETSc having a HIP backend. We present the collaborative effort between OLCF and AMD that handles the developmental work on maturing the supporting ROCm versions for the third-party workflow tools used for solving sparse matrices with respect to the Hamiltonian of the system.

## HPCQC NEEDS SCALABLE QUANTUM ERROR CORRECTION

Presenter: Marco Ghibaudi
Affiliation: Riverlane Ltd

Quantum error correction is the foundation for Fault Tolerant operations and the stepping stone towards quantum utility. Implementing a scalable solution that can integrate with existing HPC infrastructure is a challenge on its own.
Riverlane will present its novel architecture and the system engineering required to reach the utility scale.
Deltaflow combines low-latency hardware and highly tunable compilation to capture the complex requirements of Quantum Error Correction. We will present some early results on integration, challenges and key milestones required to establish a successful integration between HPC resources and Quantum computers.

## "ALL QUANTUM" TIME SYNCHRONIZATION

Presenter: Md Mehdi Hassan
Affiliation: University of Tennessee, Knoxville

Reliable time synchronization is essential for networks, sensors, and distributed computing. While classical systems like CERN's White Rabbit offer subnanosecond precision, they cannot directly synchronize quantum devices or guarantee unconditional security. We propose a fully quantum clock synchronization method using balanced HOM interference between independently generated light pulses. Numerical simulations show our method achieves sub-nanosecond precision over metropolitan-scale fiber networks and ensures robust, secure time alignment for future quantum networks and distributed computing applications.

## KRYLOV-BASED CONTINUED FRACTION SOLVER FOR GREEN'S FUNCTIONS IN DYNAMICAL MEAN FIELD THEORY

Presenter: Mariia Karabin
Affiliation: Middle Tennessee State University

We present an initial implementation of a Krylov-based continued fraction approach to compute Green's functions of strongly correlated materials on a quantum computer. The method employs variational quantum algorithms to compute the system's ground state and iteratively construct Krylov subspaces where the Green's functions can be expressed as continued fractions. As a first step, we solve the Anderson impurity model with a single bath site and evaluate the approach's ability to resolve spectral features to assess whether this method can offer accurate results with minimal bath discretization and serve as a computationally efficient alternative to conventional solvers. While current results are for the single bath case, the framework is designed to be extendable to larger bath sizes. Ongoing and future work is focused on evaluating the robustness

and scalability of the solver in more realistic settings and assessing its suitability for integration into hybrid quantum-classical simulations of correlated materials.

## IMPROVING FAULT TOLERANCE OF CAT QUBITS

Presenter: Kevin Kissell & Paul Kassebaum
Affiliation: Alice & Bob
Dissipative cat-qubits are a promising architecture for quantum processors due to their inherent quantum error correction properties. By leveraging two-photon stabilization, they achieve an exponentially reduced bit-flip error rate with increasing distance in phase-space between their basis states, with only a linear increase in phase-flip rate. The results are then improved by implementing a squeezing deformation of the cat qubit basis states, further extending the bit-flip time while minimally affecting the phase-flip rate.

## SCALABLE ARCHITECTURE FOR QUANTUM COMPUTING @ QUANTUM ART

Presenter: Amit Ben Kish
Affiliation: Quantum Art Ltd.
Quantum computation is at an inflection point. As these processors grow in scale and reliability, the need for efficient decomposition of large circuits, becomes critical.

Trapped ions-based systems form a successful platform for quantum computing, due to their high fidelity, all-to-all connectivity, and local control. However, scalability to very large qubit numbers and gate count is still a challenge. Here, we propose and demonstrate the building blocks of a holistic, scalable architecture for quantum computing.

Quantum Art's approach [1] is centered around 4 core technological pillars: multi-qubit gates [2] capable of executing the equivalent of up to 1,000 2-qubit operations in a single step; optical segmentation into independently operating cores using laser-defined optical potentials; dynamic reconfiguration of multi-core arrays for rapid entanglement distribution; and modular, high-density 2D structures within compact footprints.
This result in systems with up to x100 more gates per second, x100 more parallel operations, and a footprint up to x50 smaller. Together with our proprietary compiler [3], these technologies enable high-speed, robust execution across large-scale systems, advantageous both for fault-tolerant digital quantum computation and analog quantum simulations.

## PAULI CHECK EXTRAPOLATION FOR QUANTUM ERROR MITIGATION

Presenter: Ji Liu
Affiliation: Argonne National Lab
Pauli Check Sandwiching (PCS) is an error mitigation scheme that uses pairs of parity checks to detect errors in the payload circuit. While increasing the number of check pairs improves error

detection, it also introduces additional noise to the circuit and exponentially increases the required sampling size. To address these limitations, we propose a novel error mitigation scheme, Pauli Check Extrapolation (PCE), which integrates PCS with an extrapolation technique similar to Zero-Noise Extrapolation (ZNE). However, instead of extrapolating to the 'zero-noise' limit, as is done in ZNE, PCE extrapolates to the 'maximum check' limit—the number of check pairs theoretically required to achieve unit fidelity. In this study, we focus on applying a linear model for extrapolation and also derive a more general exponential ansatz based on the Markovian error model. We demonstrate the effectiveness of PCE by using it to mitigate errors in the shadow estimation protocol, particularly for states prepared by the variational quantum eigensolver (VQE). Our results show that this method can achieve higher fidelities than the state-of-the-art Robust Shadow (RS) estimation scheme, while significantly reducing the number of required samples by eliminating the need for a calibration procedure. We validate these findings on both fully-connected topologies and simulated IBM hardware backends.

## QCLOUDSIM: A DIGITAL TWIN OF SCALABLE QUANTUM CLOUDS

Presenter: Waylon Luo
Affiliation: Kent State University

Quantum computing has emerged as a transformative technology capable of solving complex problems beyond the reach of classical systems. The rapid development of quantum processors has led to the proliferation of cloud-based quantum computing services offered by platforms such as IBM, Google, and Amazon. These platforms introduce unique challenges in resource allocation, job scheduling, and multi-device orchestration as quantum workloads grow increasingly complex. In this work, we present a digital twin of quantum cloud infrastructures: a framework designed to model and simulate the behavior of actual quantum cloud systems. Developed in Python using SimPy discrete-event simulation library, the framework replicates key aspects of quantum cloud environments, including detailed quantum device modeling, job lifecycle management, and job fidelity. It incorporates noise-aware fidelity estimation, making it the first of its kind to simulate superconducting gate-based quantum cloud systems at an administrative level with job fidelity. We presented use cases as proof of concept, demonstrating that our quantum cloud simulation framework can serve as a digital twin of a quantum cloud and facilitate the modeling and implementation of practical systems.

## A TVD-BASED MATRIX FRAMEWORK FOR NISQ DEVICES' GATE-LEVEL NOISE ANALYSIS

Presenter: Gabrielle MacNeil
Affiliation: University of New Hampshire

Quantum computing has the potential to solve problems that surpass the capabilities of classical systems. However, they are short of fully error-correctable quantum devices. Existing quantum systems mainly fall under the category of Noisy Intermediate-Scale Quantum (NISQ) devices, with accuracy limited by gate errors, decoherence, and hardware variability due to under-explored, in-depth noise analysis.

Existing literature studies quantum errors using circuit-wide fidelity and static noise models. There is limited work available to investigate how individual gates contribute to errors during real algorithm execution, where noise behavior can vary dynamically based on gate type, context, and position. We argue the lack of gate-level resolution could hinder the development of targeted error mitigation strategies and optimization of circuit reliability at a granular level.

In this work, we propose Total Variation Distance (TVD), a statistical measure comparing observed and ideal output distributions, as a localized diagnostic tool to address the pressing need mentioned above. Different than previous studies that evaluate end-to-end circuit behavior, we employ TVD at the segment level to isolate how specific gates and combinations contribute to output degradation. The TVD values form the foundation of our gate-level noise attribution framework, offering a structured view of how noise varies across the circuit's structure and duration. By evaluating gate sequences in context, rather than relying on full-circuit error summaries or pre-characterized noise models, our method enables targeted correction or pruning of high-noise segments. To the best of our knowledge, this is the first application of TVD as a modular, segment-wise metric for noise attribution in quantum circuits.

## QUANTUM MONSOONS

Presenter: David Nicholaeff
Affiliation: New Mexico Consortium

Los Alamos National Laboratory (LANL) established the Quantum Cloud Access Project (QCAP) with the New Mexico Consortium (NMC) in order to provide a direct conduit for LANL researchers to experiment on commercial quantum systems. The first couple of years of the project primarily focused on analog quantum systems, including D-Wave's Advantage2 platform, and QuEra's Aquila quantum computer. We are now vigorously moving into fault-tolerant, error-corrected, digital systems. Concurrently, the NMC's mission is built on education and community outreach across New Mexico.

The quantum cloud, however, has proven to be a series of quantum monsoons. Building strong collaborations between vendors and researchers, providing active resources for highly motivated students and postdocs, and establishing rigorous and physically-motivated metrics and protocols for quantum benchmarking have proven to be foundational and paramount. Our solutions have ranged from heuristic user federation to ħackathons, and our applications have spanned condensed matter to nuclear application co-design.

We share lessons learned along the path, with an eye toward proofs of many-body coherence, pushing the boundaries of computational complexity with both circuit model and adiabatic model quantum computation, and most importantly, how we intend to continue strengthening research collaborations into the future.

## DATA-DRIVEN DESIGN OF HIGH-ENTROPY MATERIALS

Presenter: Corey Oses
Affiliation: Johns Hopkins University

The Entropy for Energy ( S4E ) laboratory focuses on the discovery of materials for clean and renewable energy. Specifically, the S4E lab looks to leverage the stabilizing effects of disorder to innovate clean hydrogen production, nuclear-waste immobilization, waste-heat conversion, and energy storage. The research is in line with the Materials Genome Initiative (MGI): employing high-throughput first-principles calculations and machine learning/artificial intelligence algorithms for accelerated discovery and deployment of new materials.

## BEYOND-CLASSICAL COMPUTATION IN QUANTUM SIMULATION

Presenter: Ken Robbins
Affiliation: D-Wave

Quantum computers hold the promise of solving certain problems that lie beyond the reach of conventional computers. Establishing this capability, especially for impactful and meaningful problems, remains a central challenge. Here we show that superconducting quantum annealing processors can rapidly generate samples in close agreement with solutions of the Schrödinger equation. We demonstrate area-law scaling of entanglement in the model quench dynamics of two-, three- and infinite-dimensional spin glasses, supporting the observed stretched-exponential scaling of effort for matrix-product-state approaches. We show that several leading approximate methods based on tensor networks and neural networks cannot achieve the same accuracy as the quantum annealer within a reasonable timeframe. Thus quantum annealers can answer questions of practical importance that may remain out of reach for classical computation.

## ORNL ION TRAP PROGRAM OVERVIEW

Presenter: Christopher Seck
Affiliation: Oak Ridge National Laboratory

Trapped ion quantum platforms are robust, well-controlled, and well-understood systems in the field of quantum information science (QIS) [1]. For example, these devices have been used to simulate spin chains to probe magnetic phase transitions and execute digital quantum algorithms for simulations in multiple science domains [2]. With the onset of quantum computers that can run small algorithms, domain scientists have begun testing the efficacy of these devices to deliver useful scientific results, driving strong demand for quantum computers and simulation devices. Multiple quantum computer programming stack development, algorithm development, benchmarking, simulation, computation, sensing, and general quantum computer science projects all started within the last several years at ORNL; demand for quantum resources on-site is rapidly expanding. Here, we provide an overview of the growing trapped ion QIS program and platforms at ORNL.

## LEARNING FEASIBLE QUANTUM STATES FOR QUADRATIC CONSTRAINED BINARY OPTIMIZATION PROBLEMS

Presenter: Anthony Wilkie
Affiliation: University of Tennessee-Knoxville

Quantum computing approaches excel at solving quadratic unconstrained binary optimization (QUBO) problems, however solving quadratic constrained binary optimization problems (QCBOs) is more challenging. In this work, we develop a variational approach that creates an equal superposition of quantum states that satisfy constraints in a QCBO. The method relies on flag qubits, one per constraint, to identify when a constraint is violated or not. The resulting equal superposition can be used as an initial state for quantum algorithms that solve QUBOs/ QCBOs such as Grover's search algorithm or the quantum approximate optimization algorithm (QAOA). We test the approach on sets of one and two linear inequality constraints and find that it is able to generate an equal superposition of feasible states with a 98\% AR on average. We then use the approach to generate initial states for Grover-mixer QAOA (GM-QAOA) and find that GM-QAOA with the constraint gadgets yields significantly higher probability of measuring the optimal solution than random guessing.

## SOLVING STOCHASTIC BINARY PROGRAMS WITH QUANTUM APPROXIMATE OPTIMIZATION ALGORITHM

Presenter: Anthony Wilkie
Affiliation: University of Tennessee-Knoxville

Stochastic optimization models sequential decision-making under uncertainty, where decisions int the first stage are made prior to the realization of uncertainty in the second stage. In many formulations, uncertainty is represented by discrete scenarios evolving between stages, leading to an exponential increase in problem size as the number of scenarios grows. This poses a major challenge for classical solution techniques, particularly for binary decision problems. Quantum computing offers a promising framework for tackling such problems, especially in the domain of combinatorial optimization. The Quantum Approximate Optimization Algorithm (QAOA) is one such variational quantum algorithm designed to solve discrete optimization problems on near-term quantum devices. We propose two circuit encoding strategies for solving stochastic quadratic unconstrained binary optimization (SQUBO) problems using QAOA. The first approach encodes the deterministic equivalent (DE) formulation by introducing separate decision variables for each scenario. The second employs scenario-qubits (SQ) to compactly represent uncertainty, using conditional unitaries to apply stage- and scenario-specific cost and mixing terms. We compare the performance of QAOA under both encodings in identifying optimal stage-wise decisions, and analyze the associated quantum resource requirements, including qubit counts and circuit complexity.

## CROSS PLATFORM EXECUTION OF THE QUANTUM INTERMEDIATE REPRESENTATION

Presenter: Elaine Wong
Affiliation: Oak Ridge National Laboratory
In this poster, we present updates and progress in the development of the execution engine for the quantum intermediate representation (QIR-EE) and how it helps to enable cross-platform hybrid classical-quantum computing.

## A QUANTUM COMPUTING FRAMEWORK FOR PROTEIN STRUCTURE PREDICTION ON UTILITY-LEVEL QUANTUM PROCESSORS

Presenter: Yuqi Zhang
Affiliation: Kent State University
Although quantum computing holds great promise for simulating and optimizing complex systems, existing studies have lacked a system-level framework that can be executed on real quantum devices, especially for applications in bioengineering. In this work, we present the first experimentally validated quantum computing framework for protein structure prediction that is fully executable on utility-scale superconducting quantum hardware. This framework is designed to model the three-dimensional structure of protein pocket regions during ligand docking, addressing a long-standing gap in quantum bioengineering.

Our approach integrates the Variational Quantum Eigensolver (VQE) with a vector-based tetrahedral lattice modeling strategy, a Hamiltonian construction method grounded in the physicochemical properties of proteins, and a two-stage noise-adaptive mechanism that combines parameter estimation and circuit sampling. This end-to-end hybrid quantum–classical workflow is deployed and successfully executed on IBM's 127-qubit utility-scale superconducting quantum processor "IBM-Cleveland," demonstrating its feasibility and practical effectiveness on real hardware.