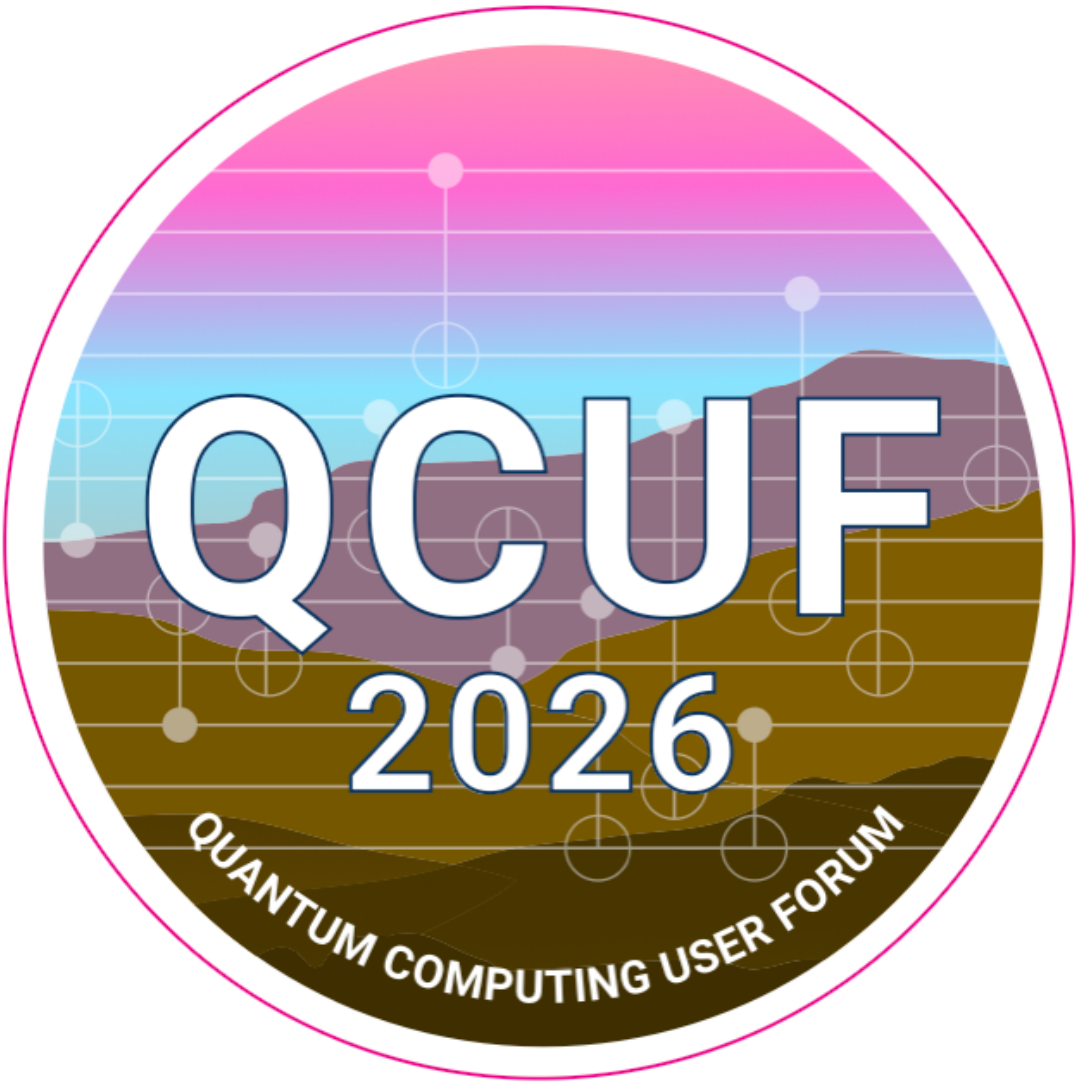


# Abstract Booklet

## Quantum Computing User Forum

## July 20–24, 2026

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

OAK RIDGE National Laboratory

## TABLE OF CONTENTS

# KEYNOTE TALKS

## HELIOS: A 98-QUBIT TRAPPED-ION QUANTUM COMPUTER

Tony Ransford—Quantinuum
Monday, July 20
11:20–12:10

Trapped-ion quantum processors provide a leading platform for quantum computation due to their high-fidelity operations and reconfigurable connectivity. Scaling system size while maintaining uniform performance, however, remains a central challenge for realizing fault-tolerant quantum computation. We present Helios, a 98-qubit trapped-ion quantum computer based on an advanced quantum charge-coupled device (QCCD) architecture designed to combine novel ion transport with high-precision quantum control. Helios utilizes hyperfine qubits encoded in $^{137}Ba^{+}$ ions and integrates multiple quantum operation regions connected through a junction transport network and rotatable storage ring, enabling flexible ion rearrangement and all-to-all connectivity. System characterization demonstrates average single-qubit gate fidelity of $2.5(1) \times 10^{-5}$, two-qubit gate fidelity of $7.9(2) \times 10^{-4}$, and state preparation and measurement (SPAM) error of $7.9(2) \times 10^{-4}$, which are maintained across operational zones. These performance levels support deep circuit execution together with mid-circuit measurement and feedforward required for quantum error correction. Beyond the component-level benchmarks, system-level benchmarks including mirror benchmarking, binary randomized benchmarking, and random circuit sampling are characterized. We show that simulating the results of the random circuit sampling is far beyond the ability of modern supercomputers. These results demonstrate that the QCCD approach enables simultaneous scaling of qubit number and operational fidelity, establishing a practical pathway toward large-scale fault-tolerant quantum computing with trapped ions. [https://arxiv.org/abs/2511.05465]

## IONQ'S PATH TO FAULT TOLERANCE

Coleman Collins—IonQ
Tuesday, July 21
11:20–12:10

IonQ is advancing trapped-ion quantum computing toward large-scale, fault-tolerant systems. This talk covers our technology and road map, connecting near-term capabilities to the longer arc toward fault-tolerant machines that complement exascale computing for scientific discovery. We discuss our current-generation systems and R&D progress on next-generation electronic qubit control systems and describe our recently published "Walking Cat" fault-tolerant architecture. We also share recent progress on scientifically and commercially relevant workloads and provide our perspective on how quantum-classical integration can bring quantum processors into high-performance computing environments.

## FROM CLOUD ACCESS TO QUANTUM ADVANTAGE: A DECADE OF OPEN SCIENTIFIC DISCOVERY AND THE ROAD TO FAULT TOLERANCE

Pedro Rivero—IBM
Wednesday, July 22
11:20–12:10

In the 10 years since quantum processors first arrived on the cloud, the field has evolved from a theoretical curiosity into an experimental tool for scientific discovery. This keynote reflects on a decade of open-access and open-source development while pivoting toward the "now" of quantum-centric supercomputing. To standardize this current paradigm, we introduce a reference architecture designed to bridge the gap between pre-fault-tolerant and early-fault-tolerant eras—working alongside the US Department of Energy in both missions. We explore the transition from proof-of-concept results to advantage benchmarks through recent milestones, including collaborative algorithmic breakthroughs with partners such as Oak Ridge National Laboratory, Cleveland Clinic, RIKEN, Q-CTRL, and BasQ. We argue that even when utilizing devices remotely, treating the quantum processing unit as a laboratory instrument is essential for discovery, and we define the best practices required for hardware-aware innovation.

## PROGRESS ON SUPERCONDUCTING QUANTUM PROCESSORS AT IQM

Kristine Rezai—IQM
Thursday, July 23
11:20–12:10

Building scalable quantum computers requires navigating a rapidly evolving landscape of scientific, engineering, and architectural challenges. This talk highlights recent progress in two superconducting qubit quantum computing architectures based on transmon qubits and flux-tunable couplers: the square-lattice IQM Crystal and the star-lattice system. I present the latest experimental results and outline the opportunities and challenges that can shape their development toward larger-scale quantum processors.

# WORKSHOPS

## MONDAY- QUANTINUUM

### Applications Implemented on Quantinuum Trapped-Ion Quantum Computers

We present recent experiments utilizing the unique capabilities of Quantinuum trapped-ion quantum computers to realize applications in materials science and quantum error correction.

In the first case, we present quantum simulations in the transverse-field Ising model and Fermi-Hubbard model at scales challenging classical methods. In the transverse-field Ising model, we study thermalization dynamics on square and triangular lattices, revealing an emergent hydrodynamics and phenomena associated to lattice frustration. The Fermi-Hubbard model is the starting point for the simulation of many strongly correlated materials, including high-temperature superconductors. We discuss recent work on the Helios and H2 trapped-ion quantum computers exploring fermionic dynamics and measuring significant superconducting pairing correlations in three regimes of the Fermi-Hubbard model, both in and away from equilibrium.

Towards error-corrected applications, we demonstrate encoding, QEC cycles, state preparation, gate benchmarking, and large-scale computation in high-rate quantum codes, using real-time quantum logic to achieve break-even performance at a quantum-advantage scale. These results show that the components of a fault-tolerant computation can be assembled and run together with better-than-physical performance, rather than demonstrated only in isolation.

### Hybrid Applications

Hybrid applications, such as quantum chemistry or quantum AI, combine HPC systems with quantum acceleration on-premises or in the cloud. These applications grow increasingly complex from a hardware and software perspective and require novel approaches to solve effectively. In this session, we will discuss state-of-the-art hybrid quantum software and its applications. Specifically, we will focus on a high-level approach for scientific workflows with Quantinuum's software stack. We will introduce concepts of workflow management systems (WMS), their benefits – such as observability, reliability, and scalability – and how to utilize quantum hardware effectively.

Concretely, users will learn how to use the Tierkreis WMS in an HPC environment, adapt existing applications to become compatible with it, and monitor the execution of an example workflow. Working along a simple example, users will gain first-hand experience in running hybrid applications.

### Real-Time Quantum-Classical Programming with Guppy

Quantinuum recently released the Helios generation of trapped-ion quantum computers, featuring 98 qubits, lower physical error rates, and faster operations than our previous systems. Alongside

these advances, Helios comes with a more powerful control system for executing dynamic quantum-classical programs. To help developers take full advantage of this capability, we introduced a new software stack centered around Guppy, our Python-based quantum programming language. This talk gives a hands-on introduction to Guppy, showing how the language makes it easy to write dynamic programs with arbitrary classical computations and measurement-driven control flow, including branching, loops, early exits, and dynamic allocation of qubits. Attendees will learn how to implement dynamic quantum algorithms as well as fault-tolerant protocols in Guppy, how to simulate them, and how to execute them on Helios.

## TUESDAY - IONQ

### Exploring the IonQ Quantum Cloud: From Hello World to Hybrid Applications

In this interactive session, we'll explore the fundamental aspects of IonQ's hardware and software ecosystem, showcase various methods for interacting with our systems, and discuss how our team utilizes these tools to pioneer novel quantum applications. The initial segment provides an overview of IonQ's trapped-ion architecture and guides participants on submitting circuits to our simulators and QPUs, using supported SDKs and our quantum cloud platform. We will then delve into specific platform capabilities and workflows, focusing on recently added functionality as well as advanced features like circuit compilation and native gates, built-in error mitigation options, and more. To conclude, we will highlight IonQ's latest advancements in quantum applications and algorithms and share an example application workflow. Throughout the workshop, attendees will have the chance to engage directly by submitting their own jobs to IonQ's simulators and hardware.

## WEDNESDAY - IBM

### Quantum Computing at Scale with IBM

Overview of current state-of-the-art of hardware capabilities and quantum computing applications, and preview of roadmap towards fault tolerance

### Tools, Methods and Best Practices on the Road to Quantum Advantage

Topics will include the spectrum from error mitigation to error correction (probabilistic error cancellation with shaded light cones, propagated noise absorption, error detection + post-selection), enhanced control over quantum computing experiments with Samplomatic and the Executor, approximate quantum compilation using tensor network methods, and real-time benchmarking for improved qubit selection

## THURSDAY - IQM

### Curvature-Augmented Imaginary-Time Evolution for Ground-State Preparation

Imaginary-time evolution (ITE) underpins a broad family of ground-state algorithms, including diffusion Monte Carlo, simulated annealing, and quantum imaginary-time evolution. Across these methods, convergence is fundamentally limited by the energy variance of the instantaneous state, a bottleneck that has remained unresolved. We introduce an augmented imaginary-time evolution framework (AITE) that replaces gradient descent on the energy landscape with a curvature-informed flow incorporating higher-order statistical moments of the instantaneous energy distribution. We prove that AITE strictly dominates standard ITE at every iteration, and since ITE arises as a special case of AITE, this acceleration propagates simultaneously across the ITE algorithmic family. We validate AITE on analytically tractable models, the hydrogen chain, and the Fermi-Hubbard model, demonstrating consistently accelerated convergence, a substantial reduction in the number of iterations required to reach a target accuracy, and superexponential convergence in imaginary time.

### A Resource-Efficient Reformulation of the Stochastic Projective Quantum Eigensolver

We improved upon the stochastic projective quantum eigensolver [1], a hybrid quantum–classical algorithm that combines stochastic sampling with the projective quantum eigensolver framework to improve convergence and reduce limitations of deterministic approaches. We propose a reformulation of the estimator that significantly lowers quantum resource requirements by simplifying the measurement process while maintaining accuracy. We benchmark results on the single Anderson impurity model and hydrogen chains to demonstrate improved efficiency, stable convergence, and good scalability. [1] M.-A. Filip, JCTP, 5964–5981, 2024.

### Integrating Applications with IQM Resonance

We will provide a generalised framework and hands-on templates to help developers and researchers leverage IQM-Resonance for seamless, application-level quantum readiness.

## LUNCH TALKS

### OLCF TODAY AND TOMORROW: HPC, AI, AND QUANTUM COMPUTING

Ashley Barker—Oak Ridge National Laboratory
Monday, July 20

This talk provides an overview of the Oak Ridge Leadership Computing Facility (OLCF), including the Quantum Computing User Program (QCUP) and OLCF's broader vision for integrating quantum computing into the future of leadership computing. Following a brief introduction to the facility, the presentation will highlight the current and future resources available through the OLCF and QCUP, including the forthcoming Lux AI and high-performance computing system; Frontier's successor, Discovery; and Oak Ridge National Laboratory's expanding portfolio of quantum platforms.

### THE QUANTUM COMPUTING ACCESS AT NERSC (QCAN) PROGRAM

Ermal Rrapaj—National Energy Research Scientific Computing Center
Tuesday, July 21

The Quantum Computing Access at NERSC (QCAN) program is a first-of-its-kind program to support research by connecting qualifying projects with our partners' quantum computers.

QCAN currently includes two paths to accessing quantum hardware: the IBM Quantum Innovation Center and QuEra Computing Collaboration. Through the National Energy Research Scientific Computing Center's (NERSC's) IBM Quantum Innovation Center, approved NERSC users may access IBM's quantum computers over the cloud for their research—an opportunity to leverage these next-generation systems for certain scientific applications. IBM's quantum computers incorporate transmon superconducting qubit technology and use tools such as Qiskit, IBM's quantum software, to make quantum computing accessible to users. Approved projects receive support, including training, to align proposals and workflows for maximum impact. Through a partnership with QuEra Computing, competitive proposals are awarded quantum computer time on the company's analog and digital neutral-atom systems, Aquila and Gemini. NERSC and QuEra staff closely engage with researchers during the project, including biweekly meetings to ensure project goals are met.

### MEET THE QUANTUM COMPUTING USER ASSISTANCE GROUP

Michael Sandoval and Jordan Winetrout—Oak Ridge National Laboratory
Wednesday, July 22

Join us for an interactive panel discussion designed to introduce our User Assistance Group and shed light on the operational experiences behind Quantum Computing User Program (QCUP) user support. The session will conclude with an open Q&A, offering you a direct opportunity to ask questions, share feedback, and collaborate on future QCUP improvements.

## CONTRIBUTED TALKS

### SIGNATURES OF TOPOLOGICAL SYMMETRIES ON A NOISY QUANTUM SIMULATOR

Christopher Lamb—Rutgers University

Topological symmetries, invertible and otherwise, play a fundamental role in the investigation of quantum field theories. Despite their ubiquitous importance across a multitude of disciplines ranging from string theory to condensed matter physics, controlled realizations of models exhibiting these symmetries in physical systems are rare. Quantum simulators based on engineered solid-state devices provide a novel alternative to conventional condensed matter systems for realizing these models.

In this work, eigenstates of impurity Hamiltonians and loop operators associated with the topological symmetries for the Ising conformal field theory in two space-time dimensions are realized on IBM's Kingston simulator. The relevant states are created on the quantum device using a hybrid quantum-classical algorithm. This algorithm is based on a variation of the quantum approximate optimization algorithm ansatz combined with the quantum natural gradient optimization method. Signatures of the topological symmetry are captured by measuring correlation functions of different qubit operators with results obtained from the quantum device in reasonable agreement with those obtained from classical computations. The current work demonstrates the viability of noisy quantum simulators as platforms for investigating low-dimensional quantum field theories with direct access to observables that are often difficult to probe in conventional condensed matter experiments.

### PARALLEL QUANTUM-HPC PIPELINE FOR CHEMISTRY APPLICATIONS

Daniel Claudino—Oak Ridge National Laboratory

In this talk, we present a quantum high-performance computing (HPC) pipeline for the fragment molecular orbital (FMO) method. The FMO method approximates the total electronic energy via a many-body expansion that keeps only few-body interactions. In our case, the calculations for each monomer and dimer are off-loaded to a quantum computer, and HPC handles the necessary orchestration and the estimation of the total energy based on each quantum computer's measurements. We illustrate this approach on helium clusters of increasing size, deploying quantum calculations to Quantum Brilliance quantum computers.

### SIMULATING LATTICE GAUGE THEORIES WITH QUBITS AND QUMODES

Victor Ale—University of Tennessee, Knoxville

Gauge theories form the foundation of the Standard Model, yet accessing their nonperturbative dynamics remains a central computational challenge. Although lattice gauge theories enable controlled discretizations, classical approaches are fundamentally limited by the sign problem and the exponential scaling of Hilbert space. I will present a hybrid quantum simulation framework that combines qumode and qubit encodings to address these limitations. In this approach, gauge fields are represented using qumodes, preserving their native infinite-dimensional structure without truncation, while matter degrees of freedom are encoded in qubits. I will construct an explicit mapping in which Hamiltonian matrix elements are evaluated in a physically motivated basis and compiled into Pauli operator expansions suitable for execution on qubit-based quantum processors. I will demonstrate the method on representative cases, including Abelian U(1) gauge theory with staggered fermions and non-Abelian SU(2) gauge theory, benchmarking spectral properties and operator expectations. The framework provides a route toward simulating real-time dynamics beyond the reach of classical computation.

## EVALUATING QPU PERFORMANCE AT SCALE—ONE YEAR LATER: NEW HARDWARE AND PERFORMANCE OF QEC PRIMITIVES

Alejandro Montanez Barrera—Forschungszentrum Julich

Quantum computers have surpassed classical simulation limits, yet noise continues to restrict their practical utility. Recently, we introduced a scalable benchmarking protocol based on the linear ramp quantum approximate optimization algorithm (LR-QAOA) to enable meaningful cross-platform comparison of quantum hardware at width (number of qubits) and depth. LR-QAOA is a deterministic variant of QAOA that probes a quantum processing unit's ability to preserve coherent signal as circuit width and depth increase, identifying when performance becomes indistinguishable from random sampling. Here, we present an updated study applying this protocol to 33 quantum processors across multiple vendors (up from 26 previously), testing up to 156 qubits and circuit depths of $p$ = 10,000. We further extend the framework to evaluate quantum error correction (QEC) primitives, focusing on mid-circuit measurement and syndrome extraction structures. We test surface code instances up to distance $d$ = 9 and color code instances up to $d$ = 11. These results establish a scalable, cross-platform methodology for benchmarking both deep-circuit performance and QEC building blocks, providing insights into the capabilities of quantum hardware.

## APPLICATIONS OF QUANTUM COMPUTING IN THE POWER GRID: ESS SITING/SIZING AND EMT SIMULATION

Phani R. V. Marthi—Oak Ridge National Laboratory

The extensive deployment of power electronics introduces spatial-temporal variability that can degrade voltage quality and operational reliability. Energy storage systems (ESS) are well suited to mitigate these impacts through fast active and reactive power support, but their value is contingent on coordinated siting and sizing. Integrated formulations that simultaneously minimize network voltage deviation, reduce substation power-flow variability, and capture installation costs often result in large-scale mixed-integer optimization problems that are computationally burdensome for classical solvers and still may not lead to the optimal solution. To address these challenges, in this talk, a two-stage hybrid quantum-classical planning framework that decouples

binary siting decisions from continuous sizing and operation will be presented. In Stage I, the siting problem is reformulated as a quadratic unconstrained binary optimization and solved using a hybrid quantum-classical workflow. In Stage II, selected locations are passed to a classical convex solver (e.g., second-order cone program) that determines optimal ESS sizes and operating set points under network constraints, ensuring physical feasibility. The effectiveness of the proposed quantum algorithm is presented through testing and evaluation of three use cases on IonQ quantum computers. In addition to the hybrid quantum-classical planning framework for ESS siting and sizing use case, electromagnetic transient simulation of power electronics building blocks using quantum state vector transformation algorithms will be presented.

## TOWARD UTILITY-SCALE ELECTRONIC STRUCTURE WITH SAMPLE-BASED QUANTUM BOOTSTRAP EMBEDDING

Joel Bierman—North Carolina State University

One of the first applications for which quantum computers are expected to demonstrate quantum advantage is finding ground states of systems relevant to materials science and molecular chemistry. However, near-term and early fault-tolerant machines will likely remain limited in the sizes of systems that can be simulated with high accuracy because of limited logical qubit counts and error rates. Quantum bootstrap embedding (QBE) is a method designed to address this issue by fragmenting a large molecular system into several smaller fragments. A quantum eigensolver then finds the ground states of each fragment in a self-consistent loop. The recently proposed sample-based diagonalization eigensolver has been demonstrated on quantum hardware for as many as 85 qubits. Its demonstrated robustness to hardware noise makes it a prime candidate as an eigensolver for experimental demonstrations of QBE on current-day hardware. In this work, we combine QBE and sample-based quantum diagonalization (SQD), denoting the resulting method as QBE-SQD, and benchmark it on ibm_pittsburgh. Our test system is a hydrogen ring that comprises eight atoms in the correlation-consistent polarized valence double-zeta basis. We find that QBE-SQD using a 43-qubit footprint per fragment often outperforms unfragmented SQD using a 67-qubit footprint by several millihartree. These results demonstrate that bootstrap embedding is a promising technique for extending the capabilities of near-term hardware.

## SHAPED CONTROL PULSES WITH BUILT-IN DYNAMICAL PROTECTION: SPACE CURVE QUANTUM CONTROL FOR DYNAMICAL DECOUPLING ON IQM DEVICES

Hisham Amer—Virginia Tech

We investigate how space curve quantum control (SCQC), a shaped-control framework that realizes arbitrary single-qubit unitaries, simultaneously exhibits dynamical decoupling (DD)–like robustness to noise, despite not being constructed as a standard DD sequence. The main advantage of unitaries implemented with SCQC over conventional approaches, which separate computation from protection, is that they use the control trajectory itself as an inherent source of noise suppression, enabling the same pulse family to function as a gate primitive, an idle-time filler, and a robust building block for digital simulation. We benchmark SCQC against Gaussian-based control and established DD families, with emphasis on dephasing and control-amplitude errors. SCQC shows strong intrinsic robustness, including in arbitrary-gate and Trotter settings in which DD-like protection is obtained without inserting separate DD layers, even outperforming traditional DD sequences in some scenarios. These findings suggest that SCQC can merge computation with error suppression into a single pulse-design paradigm without any trade-offs between high-fidelity gate implementation and qubit protection. We discuss how this perspective can be deployed on IQM devices with pulse-level access, where SCQC may provide improved idle-time management and enhanced digital quantum simulation under realistic hardware constraints.

## QUANTUM SINGULAR VALUE DECOMPOSITION AND TRANSFORM FOR SOLVING MULTIDIMENSIONAL PARTIAL DIFFERENTIAL EQUATIONS IN COMPUTATIONAL FLUID DYNAMICS AND NUCLEAR PHYSICS

Esam El-Araby—University of Kansas

The advent of quantum computing promises exponential speedup compared with existing classical methods and could alleviate computational constraints posed by multidimensional partial differential equations (PDEs) in computational fluid dynamics (CFD) and nuclear physics problems. Recent advances in quantum algorithms, software, and hardware provide a path toward realizing this goal. Quantum linear solver algorithms such as Harrow-Hassidim-Lloyd and variational quantum linear solvers have been successfully implemented to solve for canonical problems such as Hele-Shaw flow in CFD and the neutron transport equation (NTE) in nuclear physics. However, these algorithms still suffer from low scalability, high execution times, and low accuracy, particularly for problems with ill-conditioned Jacobians. In this work, we alleviate these restrictions with two new quantum singular value decomposition (QSVD) and quantum singular value transform (QSVT) PDE solvers. We validate and evaluate our algorithms by solving multidimensional PDEs for problems in CFD, such as the potential flow past a 2D cylinder and Joukowsky airfoil as well as the lattice Boltzmann method for a lid-driven cavity. Furthermore, we demonstrate the applicability of our algorithms for solving the multidimensional Boltzmann transport equation in nuclear physics by using the diffusion method for NTE in nuclear reactor engineering. Our results demonstrate higher accuracy, higher scalability, and faster execution times compared with existing solvers on noise-free and noisy quantum simulators from IBM

Quantum. We also provide analytical comparisons in terms of the temporal and spatial complexities of the two proposed QSVD and QSVT techniques.

## MANY-BODY SIMULATIONS OF INHOMOGENEOUS SUPERNOVA NEUTRINOS ON QUANTUM COMPUTERS

Zoha Laraib—University of Tennessee, Knoxville

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. Although the mean-field approximation efficiently models neutrino quantum mechanics, its accuracy to represent flavor changes remains unclear in a precise many-body treatment. To address this uncertainty, we conduct many-body simulations of neutrino quantum kinetics, focusing on supernova instabilities. The purpose of this research is to connect mean-field and many-body models on quantum computers to reveal how quantum computers can aid or hinder our understanding of the kinetic theory, addressing instabilities in entangled environments.

## DISTRIBUTED QUANTUM OPTIMIZATION FRAMEWORK FOR HIGHER-ORDER OPTIMIZATION PROBLEMS

Seongmin Kim—Oak Ridge National Laboratory

Quantum computing (QC) offers a promising way for tackling combinatorial optimization problems with exponentially large search spaces, but its practical application to real-world problems remains limited by constraints in qubit counts, restricted circuit depth, and the reliance on simplified quadratic formulations. In many scientific domains, optimization problems are inherently higher-order and densely connected, making quadratic approximations inadequate and often leading to degraded solution quality. In this work, we develop a distributed quantum optimization framework (DQOF) designed to address large-scale higher-order unconstrained binary optimization (HUBO) problems. By leveraging integration of high-performance computing with QC, DQOF can be used to execute multiple variational quantum circuits in parallel, enabling efficient exploration of complex energy landscapes. Furthermore, we introduce a clustering strategy that aggregates multiple subproblems into a single wide quantum circuit while maintaining fixed circuit depth, thereby significantly improving quantum hardware utilization. Benchmark results demonstrate that DQOF achieves superior solution quality and scalability compared with conventional quantum and classical approaches, solving HUBO problems with up to 500 variables within 170 s. These results establish DQOF as a viable pathway toward practical quantum utility in real-world scientific optimization tasks.

## TENSOR NETWORK COMPRESSION USING FLUID DYNAMICS AS A TEST BED

Dean Johnstone—Aegiq

High-performance computing systems produce extreme-scale datasets that require sampling or compression if they are to be used to their full potential. Existing data compression techniques typically exploit features such as sparsity in the data, homogeneity in the data, or a priori knowledge of which subsets of data are of most interest. Fluid dynamics data in general does not exhibit these features and so is an attractive test bed for generic compression techniques that are objective, robust, and tunable with respect to information lost due to compression. In the Oak Ridge Leadership Computing Facility Director's Discretion project Tensor Approaches for CFD Optimization (TACO), we are studying a novel quantum-inspired method based on tensor networks that meets these requirements. We will discuss a detailed and completed initial study of 1D computational fluid dynamics (CFD) data and demonstrate lossless compression. We will then outline current work to scale up this approach to higher dimension CFD problems, including 2D cuts of complex datasets, where initial results show promise, and the creation of a specific GPU-accelerated library of functions for scale-up on Frontier. We will also discuss the relation of the TACO project to the quantum CFD algorithm development being pursued in Aegiq's Quantum Applications team, including proof-of-concept demonstrations of end-to-end quantum CFD algorithms for meshes of billions of cells and portability to Aegiq's QGATE (Quantum Gate Architecture via Teleportation and Entanglement) approach to quantum computing.

## TIME PROPAGATION OF NONLINEAR DYNAMICS ON A QUANTUM PROCESSOR

Felix Motzoi—Forschungszentrum Julich

Encoding partial differential equations on quantum processors remains a significant challenge due to the nonlinear and generally non-Hermitian structure of classical dynamical systems, which necessitates additional resources and leads to substantial overhead on near-term quantum hardware. Here, we implement a hybrid quantum-classical variational framework on Oak Ridge Leadership Computing Facility hardware to address these limitations by solving both the viscous and inviscid forms of the Burgers equation. Our approach involves encoding individual terms of the partial differential equation into parameterized quantum circuits, and we explore gradient-free and sequential optimization strategies to determine circuit parameters corresponding to the field solution. To mitigate the impact of noise inherent to noisy intermediate-scale quantum devices, we incorporate error mitigation techniques such as dynamical decoupling and zero-noise extrapolation. Our results demonstrate that despite current hardware limitations, quantum processors can be used to reconstruct the time evolution of the field with high accuracy.

## FAULT-TOLERANT AND HYBRID QUANTUM LINEAR SOLVERS FOR FLUID DYNAMICS

Murali Gopalakrishnan Meena—Oak Ridge National Laboratory

Solving partial differential equations (PDEs) arising in fluid dynamics remains one of the most challenging problems in scientific computing, even as high-performance computing (HPC) enters the exascale era. Recent advances in quantum algorithms for linear systems offer a promising avenue for accelerating such computations. As the quantum hardware community moves toward fault-tolerant quantum computing alongside hybrid quantum-HPC paradigms, it is increasingly important to assess the practical applicability of both fully fault-tolerant algorithms and near-term hybrid approaches.

We present updates from our ongoing Quantum Computing User Program (QCUP) project at the Oak Ridge Leadership Computing Facility to develop, implement, and benchmark quantum algorithms for PDEs applied to canonical fluid flow problems. We examine the practical challenges of implementing the fault-tolerant Harrow-Hassidim-Lloyd algorithm on real quantum hardware and discuss algorithmic refinements aimed at improving performance. We also explore hybrid approaches—including the variational quantum linear solver and a singular value decomposition–based quantum linear solver—designed to enhance scalability on current noisy intermediate-scale quantum devices. Finally, we present preliminary demonstrations using tensor network methods for solving nonlinear PDEs, with a focus on data compression for fluid dynamics applications.

These updates highlight our QCUP project's new algorithmic advances, application milestones, and collaborations spanning national laboratories, academia, and industry over the past year.

## DQAOA FOR LARGE-SCALE PORTFOLIO OPTIMIZATION ON QUANTUM-HPC ECOSYSTEMS

Azeddine Kasmi—Citigroup

Portfolio optimization represents a significant computational challenge within quantitative finance. Recent quantum approaches have predominantly utilized quadratic unconstrained binary optimization (QUBO) or higher-order unconstrained binary optimization (HUBO) with binary Markowitz formulations to facilitate execution on gate-based hardware. Although effective for small instances, conventional quantum approximate optimization algorithm (QAOA) implementations encounter scalability and density limitations when addressing large binary portfolio problems. These limitations are particularly evident when modeling capital constraints or asset-weight discretization. Consequently, a rapid increase in qubit count and circuit depth occurs, thereby restricting the practical application of QAOA on contemporary quantum devices.

We investigate the application of the distributed QAOA (DQAOA). DQAOA enhances QAOA's scalability by effectively integrating high-performance computing with quantum computing resources. In this approach, large-scale QUBO/HUBO instances are decomposed into smaller, tractable subproblems, which are processed in parallel, and their solutions are iteratively aggregated to converge toward a global optimum.

We evaluate DQAOA on binary-encoded portfolio selection tasks, specifically targeting problems involving 50–60 assets. The performance of DQAOA is benchmarked against standard QAOA and classical optimizers, and we consider metrics such as the approximation ratio, scalability, and time to solution. This study highlights DQAOA as a promising direction for developing scalable quantum finance workflows, potentially enabling quantum-accelerated optimization for portfolio problems previously considered intractable.

## ENTANGLEMENT BENCHMARKING IN QUANTUM SIMULATIONS OF SPIN SYSTEMS

Anshumitra Baul—Oak Ridge National Laboratory

Quantum phase transitions in many-body systems give rise to highly entangled states, and understanding these correlations is crucial for characterizing quantum materials. Traditional entanglement measures such as entanglement entropy are limited to pure states and require full state tomography, making them impractical for near-term quantum devices. Therefore, we explore the positive partial transpose (PPT) criterion as an efficient and scalable entanglement witness in quantum spin models. It detects pairwise entanglement from reduced density matrices, distinguishes quantum from classical correlations, and applies to both pure and mixed states—making it ideal for studying condensed matter systems at finite temperatures. We prepare ground states of 1D spin systems using matrix product states and map them to quantum circuits via adiabatic evolution for efficient realization on quantum hardware. The PPT criterion captures clear signatures of phase transitions and is implemented on real hardware using quantum overlapping tomography. This framework provides an efficient way to completely map the two-body entanglement structure in any state, including time-evolved and thermal states, on quantum computers.

## PERTURBATIVE QUANTUM FIELD THEORY SIMULATIONS ON QUANTUM DEVICES

Herschel Chawdhry—Florida State University

Perturbative quantum field theory calculations are vital for precisely predicting the outcome of scattering processes in high-energy physics. Such calculations are computationally challenging to perform, which constrains our ability to discern discrepancies between theory and experiment to uncover early hints of new fundamental particles. In this talk, I present recent advances in quantum computing methods to tackle these calculations. The results include the first indications of a potential quantum advantage and are demonstrated on a range of quantum computers via the Quantum Computing User Program at Oak Ridge National Laboratory, including a 98-qubit ion-trap machine. In time, we anticipate this work will create new opportunities for stress-testing the Standard Model of Particle Physics and revealing potential evidence of new fundamental particles awaiting discovery.

## ROBUST CHIRAL EDGE DYNAMICS OF A KITAEV HONEYCOMB ON A TRAPPED-ION PROCESSOR

Phillip Lotshaw—Oak Ridge National Laboratory

Kitaev's honeycomb model is a paradigmatic exactly solvable system that hosts a quantum spin liquid with non-Abelian anyons and topologically protected edge modes, offering a platform for fault-tolerant quantum computation. However, real candidate Kitaev materials invariably include complex secondary interactions that obscure the realization of spin-liquid behavior and demand novel quantum computational approaches for efficient simulation. Here, we report quantum simulations of a 22-site Kitaev honeycomb lattice on a trapped-ion quantum processor, both without and with nonintegrable Heisenberg interactions that are present in real materials. We develop efficient quantum circuits for ground-state preparation, achieving high accuracy with energy errors equivalent to an effective temperature $\approx 0.2$ (in units of the Kitaev interactions), which is consistent with the experimentally relevant spin-liquid regime. Starting from these states, we apply controlled perturbations and measure time-dependent spin correlations along the system's edge. In the non-Abelian phase, we observe chiral edge dynamics consistent with a nonzero Chern number, a hallmark of topological order, which vanishes upon transition to the Abelian toric code phase. Extending to the nonintegrable Kitaev-Heisenberg model, we find that weak Heisenberg interactions preserve chiral edge dynamics, whereas stronger couplings suppress them, signaling the breakdown of topological protection. Our work demonstrates a viable route for probing dynamical signatures of topological order in quantum spin liquids using programmable quantum hardware, thereby opening new pathways for quantum simulation of strongly correlated materials.

## A CASE STUDY ON INTEGRATING QUANTUM COMPUTERS WITH HPC TECHNOLOGY AT ORNL

Claire Marvinney and Amir Shehata—Oak Ridge National Laboratory

A partnership between Oak Ridge National Laboratory (ORNL) and IQM has established a new quantum high-performance computing (HPC) test bed environment that supports a 20-qubit superconducting quantum computing (QC) system integrated with HPC infrastructure. An IQM Radiance quantum computer was delivered to ORNL in fall 2025 and was subsequently installed and handed over to ORNL in May 2026. The IQM Radiance system consists of a superconducting quantum processor unit, RF control electronics, a cryogenic cooling system, and a full software stack. The entire system is housed in an environment engineered to control vibration, temperature, and electromagnetic interference to minimize noise and reach superconductivity at 10 mK. Here, we highlight the installation journey for integrating this system at the Oak Ridge Leadership Computing Facility (OLCF). The software control stack enables low-level control with pulse schedules, calibration, and the execution of experiments. We are fully connecting this QC resource to the OLCF through segmented, monitored network paths to ORNL's workflow and scheduling infrastructure, thereby enabling hybrid jobs that orchestrate classical preconditioning on compute nodes with quantum execution and postprocessing on shared file systems.

## TOWARD STUDYING SUPERCONDUCTIVITY IN THE FERMI-HUBBARD MODEL ON RYDBERG ATOMS

Nora Bauer—IonQ

We present a method for calculating the ground-state energy of the Fermi-Hubbard model by leveraging Rydberg atom processors and sample-based quantum diagonalization (SQD). By exploiting the perturbative relationship between the Fermi-Hubbard and Heisenberg models, we can sample from the Heisenberg model as prepared on the Rydberg atom processor and use the samples to diagonalize the Fermi-Hubbard model for large $U$. We include anisotropy and next-nearest-neighbor interactions and discuss the relevant regime for quasi-superconductivity in the 1D Fermi-Hubbard model. Numerical and experimental results on the Aquila quantum processor are presented for ground-state energy calculations and the chemical potential. We find that the Heisenberg model sampling in the studied regime is sufficient to converge near the ground state for up to 56 qubits, and we see a clear advantage of Rydberg atom sampling over random sampling even with 10× more samples for diagonalization. We also present a gate-based implementation of the SQD algorithm on IBM Quantum hardware for a 56-qubit Hubbard model as a benchmark. Finally, we provide a gap analysis for studying emergent superconductivity using this method.

## Q-IRIS: ENABLING CLASSICAL-QUANTUM TASK-BASED RUNTIME HYBRID WORKFLOWS

Elaine Wong - ORNL

High-performance computing systems are rapidly evolving into heterogeneous platforms that fuse quantum processing units (QPUs) with traditional classical processing units (CPUs) and graphical processing units (GPUs). This convergence calls for runtimes capable of managing both classical and quantum workloads in a unified manner. We introduce an intelligent, task-based runtime that marries the Intelligent Runtime System (IRIS) asynchronous scheduler with a quantum programming stack through the Quantum Intermediate Representation Execution Engine (QIR-EE). Our design allows programs written in the quantum intermediate representation (QIR) to be dispatched concurrently to a variety of back-ends, including multiple quantum simulators and nascent quantum processors, enabling genuine hybrid execution on a single node. To illustrate its practicality, we partition a 4-qubit and 20-qubit circuit into three sub-circuits using quantum circuit cutting via the QCut library. Each sub-circuit is simulated independently by the QIR-EE driver within IRIS, after which a classical post-processing step merges the simulation results to estimate the original observable expectation value computation. This case study demonstrates how finer task granularity can enable the parallel execution and lower the simulation burden per quantum task while preserving overall accuracy, highlighting the feasibility of our hybrid approach. Quantum simulators and hardware backends that were used for the experiments were provided by the QCUP program.

# QSC INTRODUCTION TALKS

## QSC INTRODUCTION

Josh Cunningham—Chief Operating Officer, Quantum Science Center, Oak Ridge National Laboratory

The Quantum Science Center (QSC) is working to create the first fault-tolerant ecosystem for hybrid quantum high-performance computing (QHPC). By uniting national laboratories, academic institutions, and industry partners, the QSC seeks to develop a holistic software ecosystem that combines research in hybrid algorithms, scientific applications, QHPC architectures, and experimental validation, thereby amplifying the impact of fault-tolerant quantum computing. Its efforts will help position the United States at the forefront of quantum-accelerated computing, thus benefiting science and technology globally.

## SCIENCE APPLICATIONS

Thomas Maier—Distinguished Research Staff and Section Head, Advanced Computing Methods for Physical Sciences, Oak Ridge National Laboratory

The scientific applications team works to develop and validate simulation codes for quantum materials by implementing hybrid workflows on quantum high-performance computing systems.

## ALGORITHM DEVELOPMENT

Yan Wang—Research Scientist, Quantum Computational Science, Oak Ridge National Laboratory

The hybrid algorithms team designs workflows for the quantum high-performance computing ecosystem that target quantum simulation of model materials.

## QHPC TEST BEDS

Christopher Zimmer—QHPC Architectures Thrust Lead, Quantum Science Center, Oak Ridge National Laboratory

In the focus area of quantum high-performance computing architectures, we co-design and evaluate approaches for executing applications on hybrid computing systems.

## POSTERS

### AI FOR QUANTUM AT THE QUANTUM SCIENCE CENTER: TOWARD BUILDING A CONNECTIVE LAYER ACROSS THE QUANTUM COMPUTING STACK

Tirthankar Ghosal

Quantum computing is advancing toward utility faster than the surrounding scientific infrastructure can absorb it. The field confronts a fragmented and rapidly expanding literature, the combinatorial intractability of fault-tolerant error correction, the absence of principled evaluation standards for AI systems operating in the quantum domain, and the formidable orchestration burden of coupling nascent quantum processors with mature high-performance computing (HPC). This talk presents the AI for Quantum efforts of the Agentic Software Team at the Quantum Science Center at Oak Ridge National Laboratory, advancing a unifying thesis: that AI can serve as the connective and accelerating substrate of the quantum stack, spanning knowledge curation, scientific reasoning, and HPC-integrated execution. We present a complementary portfolio of seven efforts organized along three axes. For knowledge infrastructure, we describe Quantum RAG, a retrieval-augmented generation system grounded in domain-specific corpora, and the Quantum Computing LLMWiki, a curated, machine-readable knowledge resource. For evaluation, we introduce the Quantum Computing Leaderboard, which rigorously benchmarks frontier AI models on quantum code generation, alongside benchmark development for the IEEE Quantum Week (QCE) shared-task problems. For domain reasoning and systems, we present preliminary versions of ChatQEC and FTQEC, language-model-driven assistants for quantum and fault-tolerant error correction, and a quantum computing multi-agent framework that decomposes complex quantum-HPC workflows across specialized agents and currently demonstrates strong code generation performance. We further introduce EPIC, a multi-agent framework for HPC orchestration, data management, and predictive analytics that we are extending toward quantum-HPC co-execution. Throughout, we report design rationale, quantitative findings where available, and the characteristic failure modes that frontier models exhibit on rigorous quantum tasks. We conclude by articulating how these components compose into an integrated assistance and evaluation fabric for the quantum community and by framing the open problems that must be solved to realize trustworthy, verifiable AI for quantum science.

### ACCELERATING AND SCALING QUANTUM CIRCUIT SIMULATION USING CIRCUIT CUTTING AND FPGA ACCELERATION

Xiaokun Yang, Yunhe Feng, Xuechen Zhang, Vipin Chaudhary, and Shuai Xu

This project presents a specialized hardware generation framework for accelerating and scaling quantum circuit simulation through the integration of circuit cutting and field-programmable gate array (FPGA) acceleration. Although circuit cutting enables large-scale state vector simulation on classical hardware, it also introduces substantial reconstruction overhead that becomes a major performance bottleneck. To address this challenge, the project team proposes a parameterized hardware generation framework implemented using the Chisel hardware construction language. The framework automates the design of specialized hardware architectures optimized for massively parallel computation. The proposed approach integrates a mixed low-bitwidth

quantization scheme (QS1.f), a lightweight quantum gate library, and a deeply pipelined spatially parallel architecture to enable scalable and high-performance hardware generation.

The experimental results demonstrate that the proposed framework achieves operating frequencies of up to 900 MHz on the AMD Alveo U250 Data Center Card and delivers nanosecond-scale performance through extensive spatial parallelism while maintaining high numerical accuracy (error bound: $3.05 \times 10^{-5}$) and efficient resource utilization.

FPGA-based prototyping further validates the effectiveness of the proposed design and highlights the potential of a fabricated ASIC for high-performance quantum circuit simulation using 7 nm/16 nm TSMC technology. Although the current evaluation focuses on moderate-scale circuits, the framework establishes a scalable foundation for future quantum circuit simulation through the tight integration of circuit cutting and specialized architecture design.

## APPLICATIONS OF QUANTUM COMPUTING IN EARTH SYSTEM MODELING: REVIEW

Min Xu, Muralikrishnan Gopalakrishnan Meena, Wei Zhang, Anukesh Krishnan Kutty Ambika, Salil Mahajan, Matthew R. Norman, and Katherine Evans

Recent advances in quantum computing (QC) have motivated increasing interest in its application to Earth system modeling (ESM), particularly in weather, subseasonal-to-seasonal, and hydrology prediction. Although practical applications are at an early stage, QC could potentially accelerate simulations and improve prediction in computationally intensive ESM tasks. This review summarizes recent progress in quantum-enhanced ESM, with emphasis on four closely related areas: (1) quantum neural networks, (2) variational quantum circuits, (3) quantum kernels, and (4) hybrid quantum-classical frameworks. Existing studies indicate that QC can improve prediction or achieve performance comparable to conventional approaches while using substantially fewer trainable parameters in ML models. Such advantages have been reported in various ML applications using QC, including forecasting of river discharge, solar irradiance, air quality, wind speed, and simplified atmospheric modeling. Besides ML applications, quantum algorithms have also been investigated for radiative transfer calculations, stochastic atmospheric processes, data assimilation, and the optimization of chaotic Earth system models. However, the transition from proof-of-concept studies to operational ESM is constrained by the noise and limited qubit counts of current noisy intermediate-scale quantum hardware and by the difficulty of encoding high-dimensional geophysical data into quantum representations. Future research should focus on three key directions: (1) developing robust quantum feature maps for satellite and reanalysis datasets; (2) integrating QC into end-to-end ESM workflows, including data assimilation and ensemble generation; and (3) designing domain-specific error-control strategies to improve the practical use of QC in ESM.

## BEYOND VIBE CODING: UNIFIED MULTI-FRAMEWORK QUANTUM CODE GENERATION AND EXECUTION-AWARE EVALUATION WITH LLMS

Sumera Anjum, Xiaokun Yang, Thomas Gerrity, Junhua Ding, Kewei Sha, Weijian Zheng, Song Fu, and Yunhe Feng

Large language models (LLMs) are increasingly used for quantum program synthesis, but generated circuits may be syntactically valid and executable while failing to implement the intended quantum transformation. This work presents a unified multi-framework pipeline for quantum code generation across PennyLane, Qiskit, and OpenQASM. Starting from 7,287 unique quantum circuit problems, we construct 23,010 aligned framework-specific instances and validate cross-framework equivalence with simulation-based fidelity analysis. We adapt Qwen2.5-Coder-14B using supervised low-rank adaptation fine-tuning and further refine framework-specific adapters with Group Relative Policy Optimization (GRPO). The GRPO reward is execution-grounded, combining syntactic validity, executability, and state vector fidelity against a reference implementation. Evaluation on 1,149 held-out prompts uses a deterministic Pass@1 protocol that reports text-level similarity alongside runtime-aware semantic correctness. The results show that GRPO consistently improves execution-grounded metrics across all frameworks, increasing mean fidelity to 0.720 for PennyLane, 0.600 for Qiskit, and 0.280 for OpenQASM. OpenQASM shows the largest semantic gain, with conditional fidelity improving from 0.16 to 0.47 despite modest executability gains. We also find that BLEU and CodeBLEU decrease after GRPO while fidelity improves, demonstrating that text similarity is an unreliable proxy for quantum program correctness. The findings highlight the importance of execution-aware optimization and evaluation for building reliable LLM-based quantum software generation systems.

## DISTRIBUTED QUANTUM-ENHANCED OPTIMIZATION: A TOPOGRAPHICAL PRECONDITIONING APPROACH FOR HIGH-DIMENSIONAL SEARCH

Dominik Soós, Marc Paterno, John Stenger, and Nikos Chrisochoides

Optimization problems become fundamentally challenging as the number of variables increases. Because the volume of the search space grows exponentially, classical algorithms frequently fail to locate the global minimum of nonconvex functions. Although quantum optimization offers a potential alternative, mapping continuous problems onto near-term quantum hardware introduces severe scaling limits and barren plateaus. To bridge this gap, we propose the Distributed Quantum-Enhanced Optimization (D-QEO) framework. Instead of forcing the quantum processor to find the exact minimum, we use it simply as a topographical preconditioner. The quantum processing unit maps the landscape to locate the most promising basin of attraction, generating high-quality seed points for a classical GPU-accelerated solver to refine. To make this approach viable for utility-scale problems, we exploit the mathematical structure of separable functions. This allows us to cut a 50-qubit (i.e., $2^{50}$) global search space into independent and manageable subspaces using 5-qubit subcircuits. By executing these fragments concurrently with CUDA-Q, we completely bypass the overhead of cross-register entanglement and classical tensor knitting for separable functions. Benchmarks on the 10D Rastrigin and Ackley functions show that D-QEO prevents the exponential failure rates observed in purely classical algorithms. Furthermore, this quantum warm-start significantly reduces the number of classical Broyden-Fletcher-Goldfarb-Shanno iterations

required to converge, providing a highly practical blueprint for utilizing near-term quantum resources in complex global search.

## ERROR MITIGATION FOR QUANTUM APPLICATIONS AND SIMULATIONS BEYOND CLASSICAL LIMITS

Qedma Quantum Computing

Quantum processors are beginning to operate in regimes that challenge the best available classical simulation methods. As quantum hardware continues to improve, efficient error mitigation is becoming essential for obtaining accurate results from increasingly complex quantum circuits.

In this poster, we present recent advances in Qedma's characterization-based error mitigation software, QESEM (Quantum Error Suppression and Error Mitigation), and demonstrate how it enables high-accuracy quantum simulations on IBM Heron and Quantinuum H2 processors. We highlight recent results on quantum many-body dynamics in parameter regimes that challenge state-of-the-art tensor-network and Pauli-propagation methods on large-scale supercomputers. These experiments reveal physical behavior that cannot be reliably reproduced using current classical approaches despite extensive computational resources, thus providing an important step toward verifiable quantum advantage. We also discuss additional demonstrations spanning quantum chemistry, optimization, quantum ML, and hybrid quantum-classical workflows.

Finally, we briefly discuss recent directions combining error mitigation with quantum error correction, and we describe the role of error mitigation in the transition toward scalable fault-tolerant quantum computing.

## EXPLICIT BLOCK ENCODING OF THE DIFFERENCE-OF-GAUSSIAN OPERATOR

Jishnu Mahmud, John Winship, Tom Lash, James Ostrowski, and Rebekah Herrman

The difference of Gaussians (DoG) is a widely used operator across applications, including image processing (feature and edge detection), quantum ML, and finite-difference methods (approximations of the Laplacian of Gaussian). In this poster, we construct an explicit quantum block encoding of the DoG operator on a periodic grid, exploiting its natural probabilistic structure. The central observation is that the DoG admits a natural decomposition into two normalized Gaussian distributions, each preparable by explicit and efficient circuits, with the negation encoded using a single Pauli-$Z$ gate on a branch-indicator qubit. This decomposition enables the operator's block encoding to be directly mapped to the Linear Combination of Unitaries framework without requiring signed amplitude loading, quantum random-access memory, or any other black-box oracles. The proposed method achieves a constant subnormalization factor $\lambda = 2$ independent of the grid size $N$, the spatial dimension $D$, and the stencil width. Additionally, we show that the DoG operator is diagonalized by the discrete Fourier basis, which allows us to derive an exact closed-form expression for the block-encoding success probability in terms of the input signal's power spectrum, weighted by the operator's transfer function. Finally, we prove that the expression reduces to $O(h^4)$ scaling with respect to grid spacing $h$ as the periodic grid becomes finer. This implementation provides an explicit construction method for a tunable, wide-stencil band-pass filter whose frequency response is controlled by two Gaussian scale parameters.

## FTQC-READY AL-DQAOA FRAMEWORK FOR ENERGY MATERIALS SCIENCE

In-Saeng Suh, Seongmin Kim, DoHoon Kim, Sinchul Yeom, and Mina Yoon

High-entropy oxide (HEO) electrocatalysts offer a promising pathway toward next-generation energy technologies due to their vast compositional design space and potential for enhanced catalytic activity, stability, and selectivity. However, discovering optimal HEO compositions remains a grand challenge because of the combinatorial complexity associated with multi-element materials and competing performance objectives. In this work, we will develop a fault-tolerant quantum computing (FTQC)–ready active learning distributed quantum approximate optimization algorithm (AL-DQAOA) framework that combines quantum error correction, AI, high-performance computing (HPC), and quantum optimization to accelerate the inverse design of HEO electrocatalysts. The framework will leverage surface code–based logical QAOA on the NVIDIA Ising Decoding AI model to perform optimization on logical qubits, thereby enabling fault-tolerant quantum execution. The resulting closed-loop AI-FTQC-HPC framework will autonomously explore vast HEO composition spaces, identify high-performance electrocatalyst candidates, and establish a novel paradigm for applying FTQC to complex inverse-design problems in energy materials science.

## EXTENDING QAOA-GPT TO HIGHER-ORDER QUANTUM OPTIMIZATION PROBLEMS

Leanto Sunny, Abhinav Rijal, and George Siopsis

We trained a GPT-based model to generate adaptive derivative assembled problem tailored (ADAPT) quantum approximate optimization algorithm (QAOA) circuits for cubic higher-order unconstrained binary optimization (HUBO) spin-glass problems. Evaluated on 8- and 16-qubit HUBO instances, the model achieves average approximation ratios exceeding 0.95, closely matching classically optimized ADAPT-QAOA solutions. The generative model produces adaptive QAOA-style circuits and corresponding variational parameters in a single forward pass, eliminating the need for iterative classical optimization during inference. These results demonstrate that transformer-based generative models can generalize beyond quadratic optimization problems to higher-order Hamiltonians and more complex energy landscapes. This work highlights a scalable ML-driven pathway toward automated variational circuit design and accelerated quantum optimization in the noisy intermediate-scale quantum era.

## HIGH-ORDER CPTP INTEGRATORS FOR THE LINDBLAD EQUATION USING TENSOR NETWORKS

Peter DelMastro, Daniel Appelö, and Yingda Cheng

We propose a family of low-rank, completely positive and trace-preserving schemes for the Lindblad equation, a common model for open quantum systems. Low-rank representation is employed at two levels: the density matrix is factorized into the product of tall-skinny matrices, and the columns of these matrices are further represented using the tensor train (TT) format, also known as the matrix product state (MPS). This two-level low-rank format fits naturally into our

existing Kraus Is King scheme (arXiv:2409.08898v2) for the Lindblad equation, the underlying operations of which are arithmetic on the columns of the tall-skinny matrices. We show how these operations can be performed efficiently in the TT/MPS format, with particular emphasis on density matrix rank truncation. We conclude with extensive numerical experiments demonstrating the convergence of this scheme and its efficiency in simulating systems with up to $10^{19}$ degrees of freedom using only modest compute resources.

## HYBRID PREPARATION OF GROUND STATE FOR EARLY FAULT-TOLERANT QUANTUM HARDWARE

Aaron Seymour, Alessandro Baroni, and Antonio Coello Pérez

We present a hybrid classical-quantum workflow for ground-state preparation and energy estimation on early fault-tolerant quantum hardware, and we use the 1D Heisenberg chain as a benchmark system. The approach combines three complementary stages: (1) neural networks to estimate spectral quantities such as lowest or highest eigenvalues and gaps; (2) tensor-network ground-state approximation via the density matrix renormalization group (DMRG) and compilation of the resulting matrix product state into a hardware-compatible quantum circuit; and (3) quantum filtering to coherently refine the prepared state.

The primary state-preparation route maps DMRG-optimized matrix product states directly to quantum circuits, and the generalized unitary cluster Jastrow ansatz is explored as a variational alternative. Prepared states are subsequently refined using spectral time filtering implemented through optimized Hadamard-test circuits, reducing residual excited-state contamination and improving ground-state fidelity. This framework illustrates how classical learning, tensor-network methods, and quantum filtering can be integrated into a scalable pipeline for quantum simulation beyond 1D systems.

## LATENCY CHARACTERIZATION AND DEPENDABILITY PROFILING OF QUANTUM-CLASSICAL SYSTEMS ON AN EMBEDDED SOC

Piyush Sonawane, Andrew Kelly, and Naveed Mahmud

Hybrid quantum-classical algorithms rely on iterative data exchange between quantum circuit execution and classical postprocessing. As the quantum circuit size increases, communication and computation latencies within these exchanges become a critical system bottleneck. In this work, we present a system-level latency characterization of a hybrid quantum-classical learning workflow implemented on a cost-effective embedded system on chip (SoC). We evaluate a heterogeneous execution model in which the quantum forward pass is executed on either a local simulator or an actual quantum processing unit (QPU), while classical postprocessing functions such as loss computation, gradient estimation, and parameter updates are accelerated on a low-latency field-programmable gate array (FPGA) pipeline. We measure end-to-end iteration latency for quantum ML circuits, investigating CPU-FPGA and QPU-FPGA workflows. We also present a methodology for benchmarking the dependability of quantum-classical workflows and their execution on hybrid systems. A dependability profile for workloads and corresponding systems is developed using empirical measurements and evaluations of standard speed, scalability, and security metrics for a hybrid quantum-classical system. In the dependability profile, we define

iteration-level speed metrics—variance, jitter, worst-case latency, and stability—and a resource-aware scalability metric: circuit layer operations per second. The results demonstrate that the proposed dependability profile enables practical decision-making by linking system behavior to deployability and performance limitations. Our findings provide a performance blueprint for future design and integration of low-latency embedded systems with utility-scale quantum computing systems.

## MACHINE LEARNING–ENHANCED ADAPTIVE FIDELITY OPTIMIZATION IN W-STATE QUANTUM TELEPORTATION

Faisal Ahmed Chowdhury and Farah Ferdaus

Quantum teleportation based on W-state entanglement offers inherent robustness to particle loss but is fundamentally limited in fidelity for general quantum states. This fidelity is further degraded in noisy intermediate-scale quantum systems by gate infidelities, decoherence, and measurement errors. To address this challenge, we propose an ML-driven adaptive protocol that significantly improves teleportation fidelity. Our framework uses a neural network to predict the fidelity of each teleportation instance. Based on this prediction, an adaptive two-tier correction mechanism is employed. When predicted fidelity is low, optimization of a universal single-qubit correction gate (U3) is triggered; in high-fidelity cases, efficient Pauli gate corrections are used. This adaptive approach yields a substantial increase in average teleportation fidelity with minimal classical overhead. Specifically, it increases fidelity by 44.81%, raising the average teleportation fidelity from 0.582 to 0.843. This framework is also validated on IBM Quantum Heron R2 using a 3-qubit W-state, confirming stable performance on real hardware. This end-to-end hybrid pipeline enables reliable quantum state transfer on current noisy hardware and is a key step toward practical quantum networks. It also demonstrates how teleportation can seamlessly integrate into hybrid quantum-classical systems, making it suitable for emerging large-scale computing architectures.

## PARTITIONED-CONSTRAINT QAOA (PC-QAOA): STRUCTURAL STATE PREPARATION AND PENALTY ENFORCEMENT FOR QUANTUM OPTIMIZATION

Anthony Wilkie, Alexander DeLise, Andrew Del Real, Rebekah Herrman, and James Ostrowski

Constrained combinatorial optimization remains challenging for quantum algorithms because feasibility must be explicitly enforced, typically through penalty terms or problem-specific mixers. We introduce a partitioned-constraint quantum approximate optimization algorithm (PC-QAOA), which partitions constraints into those enforced structurally and those enforced energetically. Structural constraints are handled via feasible-state preparation and a Grover mixer that preserves feasibility, while the remaining constraints are enforced through penalties. We show that constraints with disjoint support can be prepared in parallel with little error accumulation. We identify broad classes of constraints (including cardinality, assignment, and flow conservation) that admit efficient structural enforcement and introduce a variational gadget construction that extends this approach to arbitrary low-support constraints. Across 500 instances spanning multiple constraint families, PC-QAOA substantially improves feasibility and solution quality at shallow depth relative to penalty-based QAOA, demonstrating the value of partial structural enforcement.

## PRACTICAL HPCQC INTEGRATION WITH QDMI

Lukas Burgholzer, Marcel Walter, Patrick Hopf, Matthias Traube, and Robert Wille

Quantum computers are moving into high-performance computing (HPC) centers, and the main challenge is now integration rather than pure hardware access. Many current software paths still depend on vendor-specific adapter chains between user software development kits, schedulers, and backend APIs. This pattern makes operations more complex than necessary and slows the transition from pilots to production workflows. We present a practical integration path centered on the Quantum Device Management Interface (QDMI). Using IQM superconducting systems as a hardware case study, we implement an IQM-backed QDMI layer and connect it to two software layers relevant to HPC centers working with quantum computers: Slurm-based job execution and Qiskit-facing user workflows. The key message is simple: integrating quantum hardware into HPC need not be a bespoke engineering effort for each backend. Once the software-hardware boundary is standardized, large parts of the stack become reusable across providers and deployment styles.

## QISKIT-KOKKOS: A KOKKOS-ACCELERATED, MULTI-ENGINE QUANTUM CIRCUIT SIMULATOR

Chao Lu, Murali Gopalakrishnan Meena, and Kalyan Gottiparthi

Accurate and scalable quantum circuit simulation is essential for algorithm development, hardware benchmarking, and quantum error correction (QEC) research, particularly as quantum systems advance toward fault-tolerant regimes. However, the heterogeneous landscape of modern high-performance computing infrastructure—spanning NVIDIA GPUs, AMD GPUs, and multicore CPUs—poses significant portability challenges for simulation software. We present Qiskit-Kokkos, a high-performance, performance-portable quantum circuit simulator implemented as a native Qiskit plug-in and powered by the Kokkos C++ performance portability framework.

Qiskit-Kokkos provides four integrated simulation engines through a unified Qiskit-compatible interface: (1) a state vector engine for exact simulation of ideal quantum circuits; (2) a tensor network engine for memory-efficient simulation of structured or low-entanglement circuits; (3) a noisy simulation engine for realistic device emulation for linear and nonlinear noise models; and (4) a QEC simulation engine supporting fault-tolerant protocol research. We enabled native support for dynamic circuit primitives: mid-circuit measurement, qubit reset, and classical feedforward (if-then branching), which are indispensable for simulating QEC codes, quantum teleportation, and adaptive quantum algorithms.

Qiskit-Kokkos's state vector simulator reached 40 qubits on 128 nodes (1,024 AMD graphics compute dies), a single distributed matrix product state simulation reached 65,536 qubits on 512 nodes with a bond dimension of 32, and the exact stabilizer engine reached 4,000 qubits. All runs are performance-portable across different CPU and GPU architectures.

## QUANTUM COMPUTING FOR NONLINEAR UNSTEADY VISCOUS FLOWS: A QUANTUM HOMOTOPY ANALYSIS METHOD (Q-HAM) WITH CIRCUIT RE-PREPARATION

Mohammad Mehedi Hasan Akash, Turag Dev, and Kourosh Shoele

Simulating nonlinear unsteady viscous flows is computationally intensive, which motivates the development of quantum-native algorithms that exploit superposition and entanglement to compress state representation and operator evaluation. Extending quantum solvers to nonlinear partial differential equations remains challenging because quantum linear algebra primitives cannot directly encode convective nonlinearities without exponential classical overhead. We develop a quantum homotopy analysis method (Q-HAM) solver for the unsteady 1D viscous Burgers equation, a canonical nonlinear test bed for quantum fluid simulation. The nonlinear solution is expanded as a convergent HAM correction series in which each order satisfies a linear equation driven by a Cauchy convolution product of lower-order corrections. This product is evaluated quantum mechanically using a j-register circuit prepared in superposition over Cauchy pair indices, with index-space addition executed via a Draper QFT adder, requiring only $O(\log N)$ qubits. Correction states are accumulated through a Linear Combination of Unitaries circuit with a re-preparation strategy that circumvents no-cloning constraints without intermediate measurement. Unsteady time advancement uses backward Euler discretization, modifying eigenvalues to ensure unconditional stability and eliminate zero-eigenvalue singularities of the steady problem. Implemented in Qiskit and validated against a classical reference, the method achieves geometric HAM convergence with relative errors below $10^{-8}$ within $M = 5$ correction orders per time step, with qubit count scaling as $O(\log N + \log M)$. These results establish a quantum pathway toward nonlinear unsteady flow simulation with logarithmic resource scaling.

## QUANTUM COMPUTING FOR SPATIALLY AND TEMPORALLY RESOLVED MODELS OF A WHOLE CELL

Muralikrishnan Gopalakrishnan Meena, Dileep Kishore, Jerry M. Parks, Luke Bertels, Dilip Asthagiri, Tom Beck, and Mitchel J. Doktycz

Whole-cell simulation, modeling a cell's entire functional systems over its life cycle, is an outstanding challenge in computational biology. Even the simplest living cells contain thousands of interacting metabolites and proteins spanning more than 5 and 15 orders of magnitude in space (nanometers to micrometers) and time (femtoseconds to hours), respectively. Furthermore, many of the governing physical and chemical properties remain incompletely characterized. Simulating such complex systems over a full cell cycle is computationally intractable on classical architectures, raising a central question: can quantum computing offer a viable path to whole-cell simulations that integrate molecular- and systems-level complexity?

We examine the potential of quantum computing across three hierarchical scales: atomistic-molecular modeling, metabolic and regulatory networks, and whole-cell spatial modeling. We present a complexity analysis in which we compare classical and quantum algorithms for representative biological problems and identify regimes of substantial theoretical speedup. We highlight algorithmic developments designed to leverage both near-term and fault-tolerant quantum architectures, and we discuss practical bottlenecks: data encoding overhead, system

conditioning, measurement constraints, and hybrid quantum high-performance computing integration. Together, these results sketch a road map for quantum-accelerated whole-cell modeling and the biological insights such multiscale frameworks may eventually enable.

## QUANTUM DATA REPOSITORY FOR ADVANCED QUANTUM MACHINE LEARNING APPLICATIONS

Wes Brewer and Manas Sajjan

Quantum ML (QML) is advancing rapidly for challenging problems in condensed matter physics, quantum chemistry, and materials science. Although QML advantages for classical datasets remain debated, its utility for intrinsically quantum data is uncontested. Quantum states, dynamics, and measurement outcomes naturally reside in Hilbert spaces inaccessible to conventional ML frameworks.

Despite significant advances in QML algorithms, the field faces major challenges related to reproducibility, benchmarking, and cross-platform validation. Unlike classical ML, in which benchmark datasets exist for image recognition, natural language processing, and recommendation systems, the quantum computing community lacks standardized, accessible data repositories, making it difficult to compare algorithms, evaluate performance across platforms, and establish best practices.

To address this challenge, we propose a scalable quantum data repository based on symmetry-constrained classical shadows and importance-sampling techniques, leveraging our intelligent sampler (SICKLE). These compressed representations provide efficient hardware-agnostic encodings of quantum states while preserving information relevant for downstream learning tasks. Datasets will be processed using our data readiness framework (REDI), which automates data preparation and AI-workflow assessment. Key applications include automated identification of quantum phases in strongly correlated many-body systems and data-driven discovery in quantum chemistry, including classification of electronic ground states by key chemical properties. This curated, AI-ready repository will serve as a shared community resource for developing, validating, and benchmarking QML methods across quantum computing, ML, chemistry, and materials science.

## QUANTUM-ENHANCED PIPELINE FOR ESTIMATING TRITIUM BINDING ENERGY FOR FUSION SCIENCE

Manas Sajjan, Seongmin Kim, In-Saeng Suh, and Thomas Beck

Fusion energy offers a promising pathway toward reliable, carbon-free baseload power, but its large-scale deployment depends on the sustainable production and recovery of tritium fuel. Lithium-containing molten salts such as FLiBe have emerged as attractive breeder blanket materials because they facilitate tritium transport and extraction more effectively than traditional ceramic breeders. However, predicting tritium behavior in these complex ionic liquids remains challenging because commonly used empirical force fields and density functional theory often inadequately capture many-body effects (e.g., polarization, dispersion, charge transfer, electronic correlation) that govern tritium binding and mobility. We present a quantum-enhanced computational framework to evaluate tritium binding energetics in molten-salt breeder materials.

Atomic configurations of the molten salt, with and without embedded tritium, are first encoded into compact binary representations using graph neural network–based preprocessing. An active learning–driven, fragment-based distributed quadratic unconstrained binary optimization solver (i.e., active learning distributed quantum approximate optimization algorithm) is then employed to efficiently identify low-energy configurations across large configurational spaces. The resulting workflow enables accurate estimation of tritium binding energies at the electronic ground-state level while maintaining scalability to complex molten-salt environments. These binding energetics provide a critical foundation for predictive models of tritium transport, retention, and extraction, ultimately supporting the design of next-generation breeder blanket materials for fusion energy applications.

## QUANTUM SINGULAR VALUE DECOMPOSITION

X. (Spancer) Dong, E. A. Coello Pérez, D. Bykov, and A. Georgiadou

In this work, we evaluate a quantum singular value decomposition algorithm on $4 \times 4$, $8 \times 8$, and $16 \times 16$ real-valued matrices. Specifically, we construct a loss function that penalizes the measurement bitstrings representing the off-diagonal element. We select parameterized quantum circuits based on their accuracy and use LAPACK as a reference. We study the performance using different optimization methods (gradient-less, gradient-based, and excitation-solve). We find that the excitation-solve method outperforms the other tested methods and enables us to extract more accurate singular values. Future work will improve the scalability of this code to handle larger real-valued matrices and extend it to complex-valued matrices.

## QUANTUM SINGULAR VALUE DECOMPOSITION AND TRANSFORM FOR SOLVING MULTIDIMENSIONAL PARTIAL DIFFERENTIAL EQUATIONS IN COMPUTATIONAL FLUID DYNAMICS AND NUCLEAR PHYSICS

Manu Chaudhary, Devon Bontrager, Alvir Nobel, Shima Mohaghegh, Trey Campbell, Jacob Spry, Naveed Mahmud, Pruthviraj Sadhankar, Shivansh Shrivas, Chao Lu, Kalyana Gottiparthi, Muralikrishnan Gopalakrishnan Meena, Zain Karriem, and Esam El-Araby

The advent of quantum computing promises exponential speedup compared with existing classical methods and could alleviate computational constraints posed by multidimensional partial differential equations (PDEs) in computational fluid dynamics (CFD) and nuclear physics problems. Recent advances in quantum algorithms, software, and hardware provide a path toward realizing this goal. Quantum linear solver algorithms such as Harrow-Hassidim-Lloyd and variational quantum linear solvers have been successfully implemented to solve for canonical problems such as Hele-Shaw flow in CFD and the neutron transport equation (NTE) in nuclear physics. However, these algorithms still suffer from low scalability, high execution times, and low accuracy, particularly for problems with ill-conditioned Jacobians. In this work, we alleviate these restrictions with two new quantum singular value decomposition (QSVD) and quantum singular value transform (QSVT) PDE solvers. We validate and evaluate our algorithms by solving multidimensional PDEs for problems in CFD, such as the potential flow past a 2D cylinder and Joukowsky airfoil as well as the lattice Boltzmann method for a lid-driven cavity. Furthermore, we demonstrate the applicability of our algorithms for solving the multidimensional Boltzmann transport equation in nuclear physics by using the diffusion method for NTE in nuclear reactor

engineering. Our results demonstrate higher accuracy, higher scalability, and faster execution times compared with existing solvers on noise-free and noisy quantum simulators from IBM Quantum. We also provide analytical comparisons in terms of the temporal and spatial complexities of the two proposed QSVD and QSVT techniques.

## QUANTUM SOLVER FOR PARTICLE-HOLE SYMMETRIC ANDERSON IMPURITY MODELS IN DMFT WORKFLOWS

M. Karabin, E. A. Coello Pérez, T. Sohail, D. Bykov, S. Ghosh, M. G. Meena, S. Kim, A. Shehata, I.-S. Suh, H. Terletska, and M. Eisenbach

Dynamical mean-field theory (DMFT) is a powerful framework for describing strongly correlated materials, but its practical application is often limited by the computational cost of solving the Anderson impurity model (AIM), particularly as bath size increases. In this work, we present and benchmark a hybrid quantum-classical impurity solver for particle-hole symmetric single impurity AIMs. The solver uses the variational quantum eigensolver ground-state preparation with a shallow, symmetry-preserving ansatz designed to conserve particle number and spin structure. A key feature of the approach is that the same optimized variational parameters are used to construct particle and hole excitations, enabling the impurity Green's function to be evaluated through a Krylov continued-fraction expansion without separately optimizing excited states.

We test the solver for multiple bath sizes and Hubbard interaction strengths under shot-limited conditions. Several optimization strategies, including COBYLA, Adam, and L-BFGS-B, are compared in terms of convergence, accuracy, and quantum measurement cost. We also evaluate a quantum-computed moment correction based on low-order Hamiltonian moments, which systematically improves variational energy estimates. The reconstructed density of states is benchmarked against the classical exact-diagonalization-based calculations, showing good agreement for smaller bath systems. These results establish practical accuracy and resource benchmarks for near-term quantum impurity solvers and outline a path toward their integration into one-shot and eventually self-consistent density functional theory and DMFT workflows for correlated real-material simulations.

## SIMPLIFICATION RULES FOR CONTINUOUS-TIME QUANTUM WALKS ON DYNAMIC GRAPHS

Mostafa Atallah, Daniel Dilley, Alvin Gonzales, Jishnu Mahmud, Zain H. Saleem, and Rebekah Herrman

Continuous-time quantum walks (CTQWs) on dynamic graphs realize quantum gates as sequences of time-evolving graph Hamiltonians, but naive constructions produce long sequences with redundancy. Simplification rules, which are graph rewrite rules that shorten a dynamic graph sequence while preserving the unitary it implements, are the CTQW analog of circuit identities in the gate model. We give CTQW realizations of the standard single-qubit gates, including the phase and rotation families, on two-vertex graphs and extend the framework to directed graphs through the Hermitian adjacency matrix. Building on the previously known simplification rules, we introduce new graph rewrite rules, among them a loop-passing rule by which a self-loop crosses an edge and reroutes a single-qubit phase across a coupling edge, and prove them by direct

expansion of the matrix exponentials. We demonstrate the calculus through worked circuit reductions and outline its use as a transpilation primitive for converting between the circuit model and the dynamic graph framework.

## SIMULATING NOISY QUANTUM SYSTEMS AT SCALE

Amit Jamadagni and Eugene Dumitrescu

The rapid growth in the number of qubits hosted by various quantum computing architectures demands scalable simulation techniques that incorporate noise. By efficiently compressing entanglement, tensor network–based computational routines have been successful at simulating quantum phenomena at scale. Although most of these methods involve studying noiseless quantum systems at scale (on the order of hundreds to thousands of qubits), the development of variants to study noisy phenomena has gained traction only in recent times. We introduce low-rank representations that efficiently encode noisy quantum states, and we present the impact of noise on quantum algorithms at scale.

## SOLARIS: VIRTUAL TEST STANDS FOR AUTOMATED QUANTUM PHOTONIC COMPUTING

Matthew Feldman, Zacharie Leger, and Alberto Marino

Quantum photonic computing has a translation problem. Experiments can generate nonclassical continuous variable (CV) states and tune Gaussian primitives, but the path from a measured trace to a reproducible hardware action is still mostly manual: choose what to measure, decide whether the calibration is trusted, update the model, retune the optics, and know when a regime should be abandoned. SOLARIS targets that gap as an experimental automation stack for quantum photonic hardware, with virtual test stands as the staging ground for AI-enabled operation.

The poster connects that software road map to two experimental anchors. First, spatial-mode four-wave mixing in warm rubidium generates squeezed mode pairs; probe-mode interference produces a four-mode cluster-state candidate, with matched pairs below shot noise and mismatched pairs noisy. Full multimode homodyne and phase-sensitive nullifiers are the next gate from resource evidence to validated cluster-state hardware. Second, an experimental CV compiler on a truncated SU(1,1) platform shows that squeezing improves the search itself: 5.4× better compiled precision, 3.6× faster representative convergence, and 4.55× lower mean time to solution.

SOLARIS is the bridge from these results to routine operation. It records configurations, calibrations, observables, and analysis choices. Then, it runs virtual experiments that stress nullifier pipelines, compiler traces, drift/loss/noise maps, and held-out checks before any AI loop touches hardware. AI enters only as a bounded proposal mechanism for actions the experiment can falsify: compile, retune, remeasure, or stop. The goal is not generic autonomy; it is verified, automated control for quantum photonic CV computing across present hardware and the virtual tests that precede it.

## VENDOR-NEUTRAL EXECUTION PROVENANCE FOR MULTI-BACKEND QUANTUM USER PROGRAMS: BINDING IDENTITY, CIRCUIT, CALIBRATION, AND RESULTS

Srikanth Vemula

User programs that broker access to commercial quantum hardware—including the Quantum Computing User Program at the Oak Ridge Leadership Computing Facility—face a structural reporting gap: vendor account logs do not provide experiment-level records that bind a named submitter, a transpiled circuit, a calibration state, and execution results into an offline-verifiable record. This poster presents a reference architecture using FIPS 204 ML-DSA-87 signed, append-only execution records across IBM Quantum, IonQ, Quantinuum, and IQM without modifying existing user workflows. The work maps the resulting records to NIST SP 800-53 AU controls and demonstrates applicability to reproducibility, project reporting, and audit requirements.