Position Papers for the ASCR Workshop on Basic Research Needs in Quantum Computing and Networking

https://www.orau.gov/ASCR-BRN-Quantum

Pavel Lougovski[1], Ojas Parekh[2] (co-chairs),
Joe Broz[3], Mark Byrd[4], Yanne Chembo[5],
Wibe A. de Jong[6], Eden Figueroa[7], Travis S. Humble[8],
Jeffrey Larson[9], Greg Quiroz[10], Krysta M. Svore[11] (eds.)

July 2023

[1]Amazon Web Services
[2]Sandia National Laboratories
[3]IBM
[4]Southern Illinois University
[5]University of Maryland
[6]Lawrence Berkeley National Laboratory
[7]Stony Brook University
[8]Oak Ridge National Laboratory
[9]Argonne National Laboratory
[10]Johns Hopkins University Applied Physics Laboratory
[11]Microsoft

Disclaimer

U.S. Department of Energy Points of Contact:

Kalyan Perumalla, kalyan.perumalla@science.doe.gov
Thomas Wong, thomas.wong@science.doe.gov
Margaret Lentz, margaret.lentz@science.doe.gov
Steven Lee, steven.lee@science.doe.gov
William Spotz, william.spotz@science.doe.gov
Marco Fornari, marco.fornari@science.doe.gov

DOI: https://doi.org/10.2172/2283320

Contents

Claudino, Daniel	Quantum Readiness and Robustness Through Sustainable Quantum Software Practices
Clayton, Connor	Software-Defined Quantum Networks
Coffrin, Carleton	The Potential of Analog Quantum Computing
Corcoles, Antonio	Quantum-Centric Supercomputing
Das, Poulami	Practical Fault-Tolerant Quantum Computing Needs Intelligent Architectures
Delgado, Andrea	Opportunities for quantum algorithmic research using nonlocal game models
Ding, Yufei	Tailoring Quantum Error Correction Codes to Sparse Hardware: Compiler and Software Innovations
Dodin, Ilya	Quantum computing for modeling of classical systems
Du, Shengwang	Hybrid Matter-Photon Modules for Distributed Quantum Computing and Networks
Duan, Yuhua	Quantum Computing and Networking for Energy Applications
Dumitrescu, Eugene	Integrating Application, Model, Algorithm, Compilation, and Error Correction Quantum Chasms With Type
Dutt, Avik	Physics-informed generalizable quantum software co-design
Eidenbenz, Stephen	Daisy-Chaining Annealing to Hamiltonian Simulation for Ground State and Dynamics Calculations
Englund, Dirk	Towards Distributed Quantum Computing: Developing a hybrid quantum interconnect stack
Fang, Bo	RapidQE2: Towards Rapid Error Estimation on NISQ Machines
Freericks, James	The nonequilibrium many-body problem: An application where quantum cowill show advantage over classical
Garcia, Jeannette	Making the Promise of Quantum Computing Practical
Givi, Peyman	Quantum-Enabled Computational Science and Engineering
Gokhale, Pranav	The Next Mile in Quantum Software
Gorshkov, Alexey	Speeding up quantum computers with mid-circuit measurements
Govia, G	New Methods for Assessing Progress in the Development of Quantum Systems

Killoran, Nathan	Quantum Programming 2.0: Hybrid classical-quantum programming in heterogeneous execution environments
Khodaei, Amin	Quantum Modeling and Optimization Design Environment (Q-MODE)
Kocia, Lucas	Geometric Dual of Non-Unitary Qubit Devices
Kowalski, Karol	Scaling up NISQ era simulations through robust hybrid computing
Krastanov, Stefan	Digital Twin Simulators and Electronic Design Automation for Quantum Hardware - Towards Interoperable and Automated Blueprint Models
Kumar, Prem	A Quantum Internet Beyond Dark Fiber: Coexistence of Quantum and Classical Light in the Optical Fiber Infrastructure for Scalable Quantum Networking
Kwiat, Paul	A Hybrid Free-space/Fiber Architecture for Long-distance Quantum Networks
Lai, Ying-Cheng	Prospects for interconnecting quantum networks
Lee, Joonho	Assessing the potential of quantum advantage in quantum chemistry with state-of-the-art quantum algorithms
LeBlond, Tyler	On the Need for Extensible Quantum Compilers with Verification
Li, Gushu	Computer-Aided Assessment of Early Fault-Tolerant Quantum Applications Towards Practical Quantum Advantage
Lienhard, Benjamin	The Next Bottleneck: Quantum Control for Large-Scale Quantum Processors
Lin, Meifeng	The Need for Better Evaluation of Realistic Quantum Processor Performance
Liu, Ji	Enhancing noise-awareness in quantum compilers
Liu, Junyu	Quantum algorithms for training large-scale machine-learning models
Martin, Michael	Quantum Networking at the Mesoscale: From Modular Quantum Processors to the Quantum Datacenter
Meier, Kristina	Space-based quantum networks are an essential component of future architecture for distributed quantum computers and quantum-enhanced secure communication

Mniszewski, Susan	Quantum Co-design and Integration for Science Applications
Moylett, Alexandra	Benchmarking early fault-tolerant quantum computers
Muelmenstaedt, Johannes	Now is the time to prepare for quantum simulation of classical fluid dynamics
Nagarajan, Harsha	Optimal Quantum Circuit Compilation: Algorithms and Software
Newell, Raymond	Photonic nanoemitters present a promising future for ideal quantum sources
Nisoli, Christiano	Exploring strongly correlated quantum spin systems with quantum computers
Nowrouzi, Kasra	Full-Stack Codesign: Crossing Abstraction Layer Boundaries for Algorithm- and Application-Informed Optimizations to Performance of NISQ-era Quantum Computers
Núñez-Corrales, Santiago	Quantum abstract machines without circuits: the need for higher algorithmic expressiveness
Núñez-Corrales, Santiago	Causality Discovery in Biomedicine as a Killer Application for Quantum Computing
Pallister, Sam	Quantum algorithms for dynamical systems
Patel, Tirthak	System Software Stack for Diverse Quantum Computing Technologies
Perciano, Talita	A Quantum Approach for Efficient Biomedical Data Analysis
Peters, Nicholas	Long-term cybersecurity applications enabled by quantum networks
Proctor, Timothy	Measuring progress towards useful quantum computation using application-centric benchmarks
Qian, Chen	Bring the Wisdom in Designing Internet Protocols to Quantum Networking
Qiao, Chunming	Quantum Computing and Networking: Embracing Clustered Quantum Computing (CQC) Systems for Tackling Qubit Limitations
Rahman, Akond	Growing the Science of Quality Assurance for Quantum Software Stacks
Reinhardt, Steve	Quantum Computing Requires Application-level Programming Models that Deliver Performance
Rios, Francisco	Towards a More Principled Quantum Computing Ecosystem

Saeed, Samah	ML-Assisted Noise-Aware Quantum Circuit Compilation
Saglamyurek, Erhan	Photon-matter quantum interfaces for scalable and networked quantum computers
Saglamyurek, Erhan	Spectral Engineering of Quantum Light
Senapati, Priyabrata	Visualizations for Uncertainty Analysis of Quantum Application Outputs
Shaw, Harry	Distributed Quantum Processing via Integrated Regional Quantum Networks Using Free-Space Optical Links
Shaydulin, Ruslan	Focus beyond Surface Code for Error-Corrected Quantum Advantage
Sheldon, Sarah	Advanced noise handling capabilities: from error mitigation to error correction
Silva, Wladimir	High Fidelity Noise-Tolerant State Preparation of a Heisenberg spin-1/2 Hamiltonian for the Kagome Lattice on a 16 Qubit Quantum Computer
Soh, Daniel	Network models on quantum photonic circuits
Sornborger, Andrew	Quantum Algorithms for Characterizing Experimental Quantum Systems
Stein, Samuel	Buffered Entangled Pairs for Asynchronous Quantum Networking
Subasi, Omer	An Assessment of Quantum Machine Learning Challenges and Opportunities
Suderhauf, Christoph	Fitting big data onto early quantum computers
Tebyanian, Habid	Hassle-free Extra Randomness from Quantum State's Identicalness with Untrusted Components
Thapliyal, Himanshu	Approximate Computing for Noise Resilient Resource Optimized Quantum Circuits
Tseng, Huan-Hsin	Quantum Privacy Preserving Artificial Intelligence (QPPAI)
Valivarthi, Venkata Ramana Raju	Resources for Heterogeneous Quantum Networking: Quantum Frequency Conversion
Verma, Dinesh	Quantum Network Assisted Applications
Wang, Yan	Classical Simulations of Quantum Computers
Wang, Fei	Quantum algorithm development for simulating open quantum systems with energy applications
Wu, Xiaodi	Hamiltonian-Oriented Quantum Algorithm Design and Programming
Wu, Wenji	On the Road to Quantum Network: Vision, Challenges, Insights, and Opportunities

Xie, Si	Timing Synchronization and Control Electronics for Scalable Quantum Networks
Xin, Yufeng	Hybrid Quantum Networks: Modeling and Optimization
Yan, Bin	Algorithm and hardware designs with native global quantum gates
Yeter-Aydeniz, Kubra	Measurement-based Quantum Algorithms for Near-term Quantum Hardware
Yoo, S. J. Ben	Challenges and Opportunities in Quantum Networking for Future Distributed Quantum Processing: New Perspective by Quantum Wrapper Networking
Young, Jeffrey	Enabling Multi-level Parallelism in Heterogeneous Quantum-Classical Programming Models
Young, Jeffrey	Evolving Quantum Programming Abstractions
Yu, Ruozhou	Toward Packet-switched Entanglement Distribution Networks
Yue, Meng	Quantum Algorithm, Networking, and Sensing Technologies for Future Power Grids
Zeng, William	Making quantum error mitigation practical
Zhang, Liang	Designing Quantum Routers for Quantum Internet
Zhang, Junyao	Q-NAS: Quantum Noise-Resilience Ansatz Searching for Variational Quantum Algorithm

ASCR Basic Research Needs in Quantum Computing and Networking

Sponsored by the U.S. Department of Energy, Office of Advanced Scientific Computing Research
July 11-13, 2023

Call for Position Papers: Workshop on Quantum Computing and Networking

Important Dates

- May 1, 2023: Deadline for position paper submission
- May 23, 2023: Notification of position acceptance
- July 11-13, 2023: Workshop
- Workshop Website: https://www.orau.gov/ASCR-BRN-Quantum

Submit Position Paper

Workshop Home

Motivation

On behalf of the Advanced Scientific Computing Research (ASCR) program in the US Department of Energy (DOE) Office of Science, we are organizing a workshop to identify priority research directions in quantum computing and networking to better position ASCR to realize the potential of quantum technologies in advancing DOE science applications.

The mission of the ASCR is to advance applied mathematics and computer science research; deliver the most sophisticated computational scientific applications in partnership with disciplinary science; advance computing and networking capabilities; and develop future generations of computing hardware and software tools in partnership with the research community, including U.S. industry. ASCR supports computer science and applied mathematics activities that provide the foundation for increasing the capability of the national high-performance computing ecosystem and scientific data infrastructure. ASCR encourages focus on long-term research to develop intelligent software, algorithms, and methods that anticipate future hardware challenges and opportunities as well as science needs (http://science.energy.gov/ascr/research/).

ASCR has been investing in quantum information science (QIS) since 2017. ASCR's QIS investments span a broad scope of research in quantum computing and quantum networking with investments in quantum algorithms and mathematical methods; the creation of a suite of traditional software tools and techniques including programming languages, compilers, and debugging; quantum edge computing; and quantum applications such as machine learning. ASCR is also funding quantum hardware research and quantum testbeds: two quantum computing testbeds are available at Sandia National Laboratories (SNL) and at Lawrence Berkeley National Laboratory (LBNL) to external collaborators, and two quantum internet testbeds are being developed by LBNL and by a collaboration between Oak Ridge National Laboratory (ORNL) and Los Alamos National Laboratory (LANL). More information about ASCR QIS investments can be found here: https://science.osti.gov/Initiatives/QIS.

ASCR research into quantum computing and quantum networking technologies is making rapid progress, and specialized systems are now commercially available. It is important for ASCR to understand the potential of these new and radically different technologies relative to conventional computing systems and for DOE-relevant applications. However, ASCR is not interested in exploring the underlying, specific device technologies at this workshop. This workshop will focus on the following two exploration areas:

1. **The quantum software stack and fundamental quantum computer science and algorithms research.** What elements of the quantum software stack need targeted investment in order to accelerate the development of quantum computing systems? What questions in quantum computer science should be addressed and what mathematical models should be explored in order to understand the potential of quantum computing? What research could spur new approaches to developing quantum algorithms?

2. **Quantum networking**. What lab-scale research in quantum networking would accelerate the development of quantum computers? Should larger-scale quantum networking research, such as space-based quantum communication, fall within ASCR's research priorities in QIS? What research on quantum networks will benefit multiple qubit platforms?

The workshop will be structured around a set of breakout sessions, with every attendee expected to participate actively in the discussions. Afterward, workshop attendees—from DOE National Laboratories, industry, and academia—will produce a report for ASCR that summarizes the findings made during the workshop.

Invitation

We invite community input in the form of two-page position papers that identify and discuss key challenges and opportunities in quantum computing and networking. In addition to providing an avenue for identifying workshop participants, these position papers will be used to shape the workshop agenda, identify panelists, and contribute to the workshop report. Position papers should not describe the authors' current or planned research, contain material that should not be disclosed to the public, nor should they recommend specific solutions or discuss narrowly focused research topics. Rather, they should aim to improve the community's shared understanding of the problem space, identify challenging research directions, and help to stimulate discussion.

One author of each selected submission will be invited to participate in the workshop.

By submitting a position paper, authors consent to have their position paper published publicly.

Authors are not required to have a history of funding by the ASCR Computer Science program.

Submission Guidelines

Position Paper Structure and Format

Position papers should follow the following format:

- **Title**
- **Authors** (with affiliations and email addresses)
- **Topic:** one or more of the following in the context of quantum computing and networking: applications, models, algorithms, compilation, error correction and mitigation, and codesign and integration
- **Challenge:** Identify aspects of current quantum computing and networking stacks that illustrate the limitations of state-of-the-art practice, with examples as appropriate
- **Opportunity:** Describe how the identified challenges may be addressed, whether it is through new tools and techniques, new technologies, or new groups collaborating in the codesign process
- **Assessment:** What would constitute success, and how would potential solutions be evaluated? If appropriate, metrics measuring success as well as estimates or projections of required quantum resources may be included.
- **Timeliness or maturity:** Why now? What breakthrough or change makes progress possible now where it wasn't possible before? What will be the impact of success?
- **References**

Each position paper must be no more than two pages including figures and references. The paper may include any number of authors but contact information for a single author who can represent the position paper at the workshop must be provided with the submission. There is no limit to the number of position papers that an individual or group can submit. Authors are strongly encouraged to follow the structure previously outlined. Papers should be submitted in PDF format using the designated page on the workshop website.

Areas of Emphasis

We are seeking submissions aimed at various levels of broadly scoped quantum computing and networking stacks:

- **Applications:**
  - fundamental mathematical kernels and standardized libraries,
  - new kinds of DOE science applications informed by quantum capabilities
  - assessment of realistic quantum advantages

    - tools for application performance modeling and estimating
    - application-inspired benchmarks and curated libraries of instances
    - applications of entanglement distribution networks
- **Computing and programming models:**
    - design and analysis of established and novel abstract quantum computing and programming models
    - models for hybrid quantum and classical computing
    - programming environments for expressing quantum algorithms
    - quantum network models and architectures
    - hybrid quantum and classical network design
    - models for distributed quantum computing
- **Algorithms:**
    - quantum algorithms admitting theoretical or empirical evidence of advantage for fundamental domains such as simulation, optimization, or machine learning
    - hybrid quantum and classical algorithms
    - quantum-inspired classical algorithms
    - classical algorithms and software systems to simulate quantum computers and networks, including tensor network and Monte Carlo simulations
- **Compilation:**
    - expanding the scope, utility, efficiency, and robustness of software stacks for quantum computing
    - approaches, algorithms, and software systems for circuit compilation and qubit mapping, routing, parameter optimization, and scheduling;
- **Error correction and mitigation:**
    - near-term quantum computing
    - networking applications
- **Codesign and integration across the quantum computing and networking stacks:**
    - impact of application requirements across the stack
    - impact of noise, fidelity, and gate execution time on algorithms and applications

While the program committee has identified the above topics as important areas for discussion, we welcome position papers from the community that propose additional topics of interest for discussion at the workshop.

Selection

Submissions will be reviewed by the workshop's organizing committee using criteria of overall quality, relevance, likelihood of stimulating constructive discussion, and ability to contribute to an informative workshop report. Unique positions that are well presented and emphasize potentially-transformative research directions will be given preference.

DOE ASCR Basic Research Needs in Quantum Computing and Communications Pre-Workshop Report

March 28, 2023

Organizing Committee

- Joe Broz, IBM
- Mark Byrd, Southern Illinois University
- Yanne Chembo, University of Maryland
- Wibe A. de Jong, Lawrence Berkeley National Laboratory
- Eden Figueroa, Stony Brook University
- Travis Humble, Oak Ridge National Laboratory
- Jeffrey Larson, Argonne National Laboratory
- Pavel Lougovski, Amazon Web Services
- Ojas Parekh, Sandia National Laboratories
- Gregory Quiroz, Johns Hopkins University Applied Physics Laboratory
- Krysta Svore, Microsoft

DOE Points of Contact

- Thomas Wong, DOE, Advanced Scientific Computing Research
- Hal Finkel, DOE, Advanced Scientific Computing Research
- Ceren Susut-Bennett, DOE, Advanced Scientific Computing Research

1 Introduction

Quantum computers are alone in their theoretical power: they are the only known devices expected to violate the strong Church–Turing thesis that all other computing devices are equivalent up to polynomial-factor overheads [1]. There are numerous challenges to creating high-quality quantum computers that are capable of realizing their theoretical power. These include the need for new algorithms (Section 4), hardware and software infrastructure including compilation (Section 5), and error correction (Section 6).

The purpose of this workshop is to bring together quantum experts from DOE and the broader quantum information science research community to (i) understand the state of the art in models, applications, algorithms, compilers, and other cross-cutting software stack components for quantum computing and networking; (ii) identify gaps in methods and tools for developing, implementing, validating, and analyzing the performance of quantum algorithms and applications; (iii) discuss pathways to accessing realistic quantum advantage in applications informed by the needs of the DOE Office of Science; (iv) assess future research needs of the scientific community; and (v) sketch out emerging research directions. In addition, the workshop is targeting methods of quantum error detection, suppression, and mitigation throughout the quantum computing and networking stack (at compilation, runtime, detection, and postprocessing) as well as noise-resilient algorithms and application.

In order to fully leverage current and future quantum computing and networking hardware in applications, it is important to research and develop a resilient software stack on top of the hardware. The OSI stack model developed for connected multilayered classical computing systems, including networks, can serve as a guiding principle for developing a quantum software stack. Following this model, the workshop will be centered on four main and two cross-cutting themes, following the components of a typical software stack. The main themes are (i) applications, (ii) models of quantum computers and networks, (iii) algorithms, and (iv) compilers. The cross-cutting themes are (i) error correction and mitigation across the stack and (ii) integration of the main stack components.

2 Applications

2.1 Quantum computing

Over the past decade, quantum computing applications to scientific computing problems have been an area of active research. Many math kernels underpinning scientific computation, such as eigenvalues and singular values of a matrix [2], linear systems of equations [3–5], linear and nonlinear differential equations [6–8], convex and nonconvex optimization [9–12], and Markov chain Monte Carlo [13], have been examined for quantum speedups in asymptotic and NISQ settings. These primitives opened an avenue for quantum machine learning (QML) and quantum simulation algorithms with applications ranging from strongly correlated materials [14] and nuclear structure [15] to fusion [16] and high-energy physics [17, 18].

End-to-end quantum speedup (also known as quantum advantage) is a major motivation for building quantum computers and, consequentially, is a sought-after property in applications. Identifying fundamental math kernels that may admit quantum advantage and quantifying bottlenecks, for example, classical-quantum input/output or access to quantum random access memory, remains a challenge. In many instances, the best-known classical algorithms that are used to evaluate quantum advantage are not optimal and hence further complicate benchmarking. Thus, developing application-inspired benchmarks, including suites of curated problem instances for math kernels, is critical for evaluating the progress in quantum algorithm design.

Beyond the search for asymptotic quantum advantage (for large instance sizes), tools for estimating and optimizing quantum resources for fixed problem instance sizes are indispensable for quantum application developers. Coupled with the development of standardized quantum libraries for math primitives analogous to classical libraries such as BLAS, this will enable domain science experts to experiment and create new quantum simulation algorithms and applications with performance guarantees.

2.2 Quantum network applications

Quantum networks distributing entanglement over intra- or intercity distances may be used in conjunction with the current internet to provide more security and privacy. Additionally, they may provide the internet with new capabilities such as distributed quantum computation, teleportation-based communication, and quantum-based metrology. Since quantum internet technologies are still nascent, only a limited set of envisioned applications have received recent attention and have developed scientific maturity. Among them are (i) secure communication, where through the use of quantum key distribution [19], quantum secret sharing [20], and quantum secure direct communication [21] network hacking attempts can be detected; (ii) blind quantum computing, allowing network users to interact with a quantum computer without revealing sensitive information to other parties [22]; (iii) clock synchronization, where a network of remotely located quantum clocks distributes high-precision time data [23]; (iv) long baseline telescopes [24], where quantum communication techniques can be applied to extend the baseline of telescopes, enhancing the optical information gathered; and (v) distributed quantum computing [25], where quantum links are used to connect (entangle) multiple quantum processors together, enhancing their computational capabilities.

3 Models of Quantum Computing and Quantum Networking

Models for quantum computing and quantum networking share a common goal of providing logical abstractions that define structure and behavior. For example, purely theoretical models have advanced to abstract machine models that are used to design new software stacks and protocol layers [26–28]. Key concerns motivating this topic are how to incorporate the fundamental properties of quantum information while adhering to architectural decisions that bound system performance.

3.1 Quantum computing

The design and analysis of architectural abstractions of quantum computers emphasize the models by which to represent and reason about quantum computation. These models includes digital, measurement-based, and analog models of quantum computation as well as their implementations across a diversity of quantum technologies for multiple application purposes [29–31]. In addition, these models must account for interactions between different computational paradigms, including the quantum-classical or so-called hybrid model of computation. For example, interactions between the quantum program and the quantum hardware layer are often managed by conventional control electronics limited by both bandwidth and memory as well as other resources.

Quantum computing models and their associated architectures play an important role in guiding expectations for performance as well as in guiding the programming and execution of implementations in physical systems [32, 33]. For example, the development of hybrid computing models, which merge quantum and classical algorithms into single workflows, depends on the model of computation as well as the interfaces for controlling the flow of information [34]. An emerging challenge is to understand how quantum computers may integrate with high-performance computing systems, which are themselves highly tuned to application performance and technology constraints [35–37]. Similar concerns carry over into the development of programming environments for expressing quantum algorithms through the dependence on data types, data structures, compilation tools, and other interfaces [38].

3.2 Quantum networks

Quantum network modeling develops a framework for understanding quantum state preparation, propagation, and detection. Quantum state preparation is essential to understanding sources of quantum states, including the optimization of metrics such as power and brilliance. Quantum state propagation accounts for transmission impairments such as losses, birefringence, dispersion, diffusion, and random noise. Analysis of quantum state detection permits calculating the expected values corresponding to the various quantum observables. Since communication between these quantum components must be set up in advance, elementary quantum communications networks have a strong resemblance to circuit-switched networks. Lessons learned from the development of circuit-switched and cellular networks suggest well-defined protocols delineated in three different planes: data, control, and management. The data plane carries the quantum and classical information flowing on the network and is responsible for transmission reliability and security. The control plane is responsible for setting up the initial mechanism for establishing entanglement distribution between nodes. The management plane is responsible for managing any errors and infidelities that happen during the entanglement distribution.

Architectures for quantum networks provide logical models for these planes as well as protocols for the control of network resources [39–43]. Nodes within these networks may consist of quantum computers and other quantum information processing platforms [23, 44]. The development of models for sources and detectors is essential to ensuring accurate estimates of network performance [45, 46]. The necessity of classical signals in quantum networks obeys a twofold objective. The first one is that most quantum communication or computation protocols require the exchange of classical bits; the second objective is associated with the fact that the operation of quantum networks requires the exchange of classical signals. Designs of hybrid networks merge the use of quantum and classical communication and memory resources and provide frameworks in which to define new protocols and evaluate their performance [47–51]. Hybrid network design is a significant consideration in the understanding of distributed quantum computing because of the concurrent dependence on message passing and entanglement sharing [52–54].

4 Algorithms

Carefully crafted quantum algorithms are an essential ingredient in realizing advantages over classical computing and algorithms. Their potential for speedups over best-known classical counterparts has driven interest in quantum computing, starting especially with Shor's prime factorization algorithm [55] and Grover's unstructured search algorithm [56]. Although a variety of quantum algorithms are currently known [57], there is generally a paucity of known rigorous superpolynomial quantum advantages [58–61], especially for fundamental problems and kernels to empower DOE science applications. Quantum algorithms admitting theoretical or empirical evidence of advantage for fundamental domains such as simulation, optimization, and machine learning are likely to impact a broad spectrum of science applications. We lack a deep understanding of the features of quantum mechanics that give rise to quantum advantages, perhaps in part because we do not have mature fault-tolerant quantum computers to help empirically vet algorithmic ideas. Focusing on higher-level problems and kernels enables quantum algorithm architects to impact a variety of potential applications. Discovering settings in which substantial quantum advantages are possible or provable remains a foremost goal, and this may benefit from considering resources beyond execution time, such as space or memory, quality or accuracy of the solution, the number of queries made to a data-access oracle, or the number of samples drawn from a distribution.

Quantum advantages may take many forms; and hybrid algorithms, through which quantum and classical computing may complement one another, are a viable strategy for leveraging near-term quantum computing. Variational algorithms [62] for quantum chemistry and discrete optimization applications have grown popular, although they have limitations [62–64] and the advantages they may provide are unclear. Novel quantum-classical computing schemes remain a research challenge.

A related front is quantum-inspired classical algorithms and mathematical proofs. Taking a quantum-mechanical perspective can inspire purely classical advances, including a recent body of work on classical algorithms for machine learning and linear algebra obtained by "dequantizing" quantum algorithms [65, 66]. Dequantization is important, not only as a source of novel classical algorithms, but also for understanding the critical features of quantum advantages. However, care must be taken to ensure that dequantizied algorithms employ realistic classical models and assumptions [67, 68]. New types of quantum-inspired classical algorithms and analysis are valuable, as are fair and rigorous models for "apples-to-apples" comparison of classical and quantum algorithms.

Classically, our attempts to solve NP-hard problems have borne new algorithmic ideas, including approximation and parameterized algorithms [69, 70] and heuristics for optimization problems [71]. The complexity class QMA is the natural quantum counterpart of NP; and physically motivated QMA-hard problems, such as computing properties of extremal eigenstates of local Hamiltonians, continue to draw attention in both the physics and the computer science [72]. The natural analogy between classical constraint satisfaction problems (CSPs) and quantum local Hamiltonians is underexplored. Recent work has unearthed new connections between the approximation of CSPs and local Hamiltonians [73–78]. A prime research avenue is further generalizing intuition and techniques for classical CSPs and discrete optimization problems to bear on physically motivated quantum problems.

Empirical prototyping and evaluation of quantum algorithms and protocols are a challenge, since emerging quantum computers are generally noisy and limited in size. Classical simulation of quantum computers and networks remains a viable approach, and classical algorithms and software systems to push the boundaries of simulation are critical. In addition to further advancing established techniques such as tensor network [79–81] and Monte Carlo [82–84] simulations, new ideas to accelerate simulation, perhaps for specialized settings where further assumptions are possible, warrant further exploration and development.

5 Compilation

The ability to express and efficiently execute complex quantum algorithms on quantum computing hardware is critical to enable the broad adoption of quantum computing by the scientific computing community [85]. So far, programming quantum algorithms have been mostly limited to low-level gates and qubit resources or libraries specific to a single scientific domain and a specific programming model (digital/gate-based versus Hamiltonian or analog versus photonic; see Section 3) and specific quantum hardware technology. There are analogs in classical computing where CPU- or FPGA-based computing technologies require distinctly different programming and compiling paradigms.

In order to achieve the efficient execution of algorithms, compilation tools should be closely tied with programming languages and hardware execution languages; and compilers capable of adapting to emerging trends in quantum programming are an important consideration for software development [86]. Quantum compiling is a hybrid process incorporating techniques from compilers for classical programming languages, transforming high-level language to assembly language, and hardware synthesis by hardware description language, where functions are automatically synthesized into customized hardware. This requires quantum programming methods and compiling techniques that are commensurate with today's sophisticated classical approaches. Compilers and their intermediate representations of hybrid programs are important for reasoning and optimizing execution in mixed machine models [87–90]. These include languages and compilers that extend conventional tools, including C/C++ and LLVM, to quantum settings [91–93], as well as executable languages that operate close to the hardware [94–98]. Domain-specific tools have also become essential for translating and compiling existing workflows to quantum technology targets [99–101].

Considerable efforts have been expended in developing compilers and compiling strategies to transform (also known as transpile) or synthesize a quantum algorithm in a series of digital or gate-based quantum operations, commonly referred to as quantum circuits, to optimize/minimize the number of operations that need to be executed on near-term devices [89, 102–112]. These include efforts to map/route circuits based on the limited connectivity of qubits in quantum hardware architectures [113–117] and to break traditional abstraction layers [102]. Additional optimizations at levels lower than circuits and gates can be achieved through the optimization of hardware pulses that define a gate operation [118]. Tools for tuning collective gate operations are also important for optimizing circuit performance [119, 120]. Constraints in the number of qubits available for quantum simulation have led to research approaches focused on cutting larger circuits to use smaller devices [121, 122].

The number of qubits in quantum hardware is growing at a rapid pace. At the same time, the connectivity between qubits is changing, with quantum computing chips being loosely coupled [123, 124]. The community is also continuing development of new types of qubit operations, such as gates between 3 or more qubits, while exploring the use of qutrits [125, 126] and qudits [123] that allow more information to be encoded in a single quantum processing unit. In order to take full advantage of these new hardware advances, new approaches will need to be developed, based on hierarchical approaches and the exploitation of common patterns and functions, to efficiently optimize the circuits.

Noise within quantum hardware leads to errors in quantum computation. Methods for mitigating errors have gained widespread acceptance, and compilation tools need to manage the encoding of these methods directly into the program [127]. While there is considerable research in error correcting codes [128–132], the development of compilers that tailor quantum error correction to specific scientific application has been limited [133]. A broader discussion on error correction and mitigation approaches can be found in Section 6.

6 Error correction and mitigation

Errors have always been one of the greatest obstacles to realizing reliable and scalable quantum information processing. Quantum systems are subject to interactions with their environments and systematic errors that ultimately degrade qubit coherence and gate performance. To achieve long relative decoherence times in accordance with the DiVincenzo criteria [134], one must manage noise in quantum systems. Typical approaches for noise management seek to correct, avoid, suppress, or mitigate errors [135, 136]. Quantum error correction utilizes logical encodings of physical qubits to detect and correct errors [1, 137–139]. As a special class of passive quantum error-correcting codes, decoherence-free subspaces and noiseless subsystems exploit symmetries in the system-environment interaction Hamiltonian to avoid noise [140, 141]. In contrast, error suppression approaches, commonly referred to as dynamical decoupling, leverage fast and strong pulses to effectively average out noise [136, 142], while error mitigation involves amplifying the noise systematically and extrapolating to the so-called zero-noise limit [143–148]. Each class of protocols offers a form of protection against the accumulation of errors that strongly depends on the underlying noise model. For such methods to be truly effective, therefore, one must have a sufficiently accurate model that can be used to inform the selection and tailoring of the protocol. Moreover, to reduce the potential resource overheads of employing a single approach alone, practical error protection will likely employ multiple techniques working together to achieve fault tolerance in both quantum computing and networking domains. The latter poses its own unique challenges, where error protection and entanglement purification are required [39].

Quantum error correction offers a potential route toward scalable quantum computing. Provided that error thresholds can be met and the noise is sufficiently uncorrelated in time and space, the set of addressable errors that can be detected and corrected by an error correcting code increases with the number of qubits. Unfortunately, such fault-tolerant error thresholds have yet to be achieved in currently available hardware. Furthermore, real hardware noise is distinct from the underlying assumptions of error correction in that it can be correlated in space and time and possess attributes of nonstationarity [149–162]. Approximate quantum error correction, where error subspaces are not fully preserved, may offer a potential pathway toward achieving reliable quantum computation in the near term without requiring stringent fault-tolerant thresholds to be met [163–165]. Addressing the increased complexity of noise models likely requires the adaptation of error correction codes to specific noise environments. In recent years it has been shown that tailoring codes to specific noise models can enable reductions in quantum error correction overhead [166]. The potential benefit of noise-tailored codes is complicated, however, by the need for more accurate noise models. Consequently, there is a need for scalable noise characterization schemes via gate set tomography [167, 168], randomized benchmarking [169, 170], quantum noise spectroscopy [171], or some combination of techniques that can yield sufficiently accurate noise models of quantum processors of increasing size.

The advantages of an accurate noise model extend beyond quantum error correction. Knowledge of the underlying noise processes enables the tailoring of error suppression [172–175], avoidance [176–179], and error mitigation [180–183] protocols. Targeted protocols can be designed to address specific noise sources, and when combined with alternative methods, could yield an overall reduction in error accumulation and resource overhead. Instances of hybrid error protection schemes have been experimentally tested in recent years in the context of quantum error correction [129, 184, 185] and quantum algorithms [186–188]. Although promising results exist, there is little analytical insight into why certain methods perform better than others. Results remain empirical, and approaches are heuristic rather than based on the characteristics of the system noise model or application.

7 Integration of stack components and other cross-cutting topics

As quantum software and hardware increase in sophistication and complexity, so does the need for techniques and tools for analyzing the quantum stack. The evaluation of individual stack components and their ultimate integration can provide key insights into the propagation of errors [189–192], algorithm performance [81, 193–198], benchmarking [199–201], and the estimation of classical and quantum resources [202–209]. When combined, this knowledge can facilitate the targeted honing of specific quantum technology stacks or enable cross-platform comparisons and the identification of broader principles for scalable quantum computing and networking architectures.

Tools for assessing error propagation are necessary for connecting application requirements to hardware specifications. Applications define allowable error thresholds in problem accuracy. How these thresholds propagate through the stack to hardware requirements such as gate fidelities, topology specifications, and decoherence rates is not well understood. Tools designed to bridge these gaps could shed light on when quantum technology is amenable to certain application domains. Furthermore, they could help inform the design of future hardware for specific applications. Equally relevant is the bottom-up analysis of error propagation. Given a particular hardware, one could inquire about the expected application performance [189, 191, 192]. This analysis could yield information regarding the potential utility of a hardware platform for a particular application.

Evaluations of algorithmic performance through simulation and benchmarking are necessary for the verification and validation of quantum algorithms and application development. In particular, there is a need for scalable modeling and simulation tools for noisy circuit execution. These include specialized models for propagating quantum states under gate operations as well as probabilistic methods for simulating experimental observables [81, 193–198]. Crucially, the efficacy of noisy circuit simulation relies on sufficiently accurate noise models of real hardware to draw reasonable conclusions about predicted hardware performance. The development of noise models requires scalable characterization protocols that offer a sufficient and efficient characterization of multiqubit processors [167, 170, 171, 210]. In principle, such protocols could enable the active monitoring of hardware performance—which in turn could inform future simulations—and quantum program performance [211, 212]. Simulations of quantum algorithms and applications must be complemented by standardized benchmarks for evaluating and comparing quantum processors. Furthermore, suitable benchmarks must be developed for testing the integration of software and hardware to evaluate performance trends in quantum computing systems [199–201].

Resource estimation has persisted as a key component of assessing the feasibility of practical quantum algorithm implementation. While many tools exist, open questions remain about the accurate estimation of classical resource requirements in addition to their quantum counterparts [205, 213]. For example, at the physical layer, how does one account for the complexity of the control waveforms in resource estimates? Or from the classical computing perspective, how does one assess the possible energy requirements for executing a quantum algorithm on hardware [209]? In principle, resource estimation tools should possess the versatility to adapt to distinct software stacks applicable in the domains of both quantum computing and networking. The estimation of resources within specific stack layers and through the stack from the hardware to the application could be advantageous for optimizing resources across the stack. The strong interplay between stack layers poses many challenges for the optimization of classical and quantum resources that, if addressed could further aid in the evaluation and design of future quantum systems.

Quantum Internet Stack.
The quantum internet will require its own stack abstraction, incorporated with the elements of the classical internet stack, allowing for its control and operation. In this quantum internet concept, on-demand quantum operations within a quantum network, such as qubit generation or Bell state measurements, will require the precise time synchronization and remote control of classical devices over ancillary classical network infrastructure. This classical network will oversee the flow of control commands responsible for the execution of quantum network operations.

In classical networks, the network functionality is based on a layered software construct. This so-called network stack is an implementation of a suite of communication protocols that enables end-to-end communication between various connected systems. The complexity in orchestrating the tasks required to create a quantum-enabled internet concept calls for a new quantum-aware control paradigm, accommodating quantum-enabling and quantum operations simultaneously. The components and functionality of this hybrid stack can be separated into three stack protocol sets: (a) the quantum stack, controlling interconnected quantum devices performing entanglement generation and distribution, (b) the quantum-enabling stack, enabling the interconnection of classical devices in a network specifically built for maintaining the quantum coherence, and (c) the classical stack, enabling the interconnection of classical devices in a network specifically built for maintaining the quantum coherence, as well as a quantum-aware control plane, which is a set of dedicated protocols within the classical and the quantum-enabling stacks required to control and orchestrate quantum network operations.

References

[1] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information.* Cambridge University Press, 2011. DOI: 10.1017/cbo9780511976667.

[2] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe. "Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics". *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing.* 2019, pp. 193–204. DOI: 10.1145/3313276.3316366.

[3] A. W. Harrow, A. Hassidim, and S. Lloyd. "Quantum algorithm for linear systems of equations". *Physical Review Letters* 103.15, 2009, p. 150502. DOI: 10.1103/PhysRevLett.103.150502.

[4] D. W. Berry, A. M. Childs, A. Ostrander, and G. Wang. "Quantum algorithm for linear differential equations with exponentially improved dependence on precision". *Communications in Mathematical Physics* 356, 2017, pp. 1057–1081. DOI: 10.1007/s00220-017-3002-y.

[5] L. Wossnig, Z. Zhao, and A. Prakash. "Quantum linear system algorithm for dense matrices". *Physical Review Letters* 120.5, 2018, p. 050502. DOI: 10.1103/PhysRevLett.120.050502.

[6] A. Montanaro and S. Pallister. "Quantum algorithms and the finite element method". *Physical Review A* 93.3, 2016, p. 032324. DOI: 10.1103/PhysRevA.93.032324.

[7] J.-P. Liu, H. Ø. Kolden, H. K. Krovi, N. F. Loureiro, K. Trivisa, and A. M. Childs. "Efficient quantum algorithm for dissipative nonlinear differential equations". *Proceedings of the National Academy of Sciences* 118.35, 2021, e2026805118. DOI: 10.1073/pnas.2026805118.

[8] A. M. Childs, J.-P. Liu, and A. Ostrander. "High-precision quantum algorithms for partial differential equations". *Quantum* 5, 2021, p. 574. DOI: 10.22331/q-2021-11-10-574.

[9] F. G. Brandão, A. Kalev, T. Li, C. Y.-Y. Lin, K. M. Svore, and X. Wu. "Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning". 2017. DOI: 10.48550/arXiv.1710.02581.

[10] S. Chakrabarti, A. M. Childs, T. Li, and X. Wu. "Quantum algorithms and lower bounds for convex optimization". *Quantum* 4, 2020, p. 221. DOI: 10.22331/q-2020-01-13-221.

[11] J. Van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf. "Quantum SDP-solvers: Better upper and lower bounds". *Quantum* 4, 2020, p. 230. DOI: 10.1109/FOCS.2017.44.

[12] A. Montanaro. "Quantum speedup of branch-and-bound algorithms". *Physical Review Research* 2.1, 2020, p. 013056. DOI: 10.1103/PhysRevResearch.2.013056.

[13] A. Montanaro. "Quantum speedup of Monte Carlo methods". *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 471.2181, 2015, p. 20150301. DOI: 10.1098/rspa.2015.0301.

[14] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven, and G. K.-L. Chan. "Low-depth quantum simulation of materials". *Physical Review X* 8.1, 2018, p. 011044. DOI: 10.1103/PhysRevX.8.011044.

[15] I. C. Cloët, M. R. Dietrich, J. Arrington, A. Bazavov, M. Bishof, A. Freese, A. V. Gorshkov, A. Grassellino, K. Hafidi, Z. Jacob, et al. "Opportunities for nuclear physics & quantum information science". 2019. DOI: 10.48550/arXiv.1903.05453.

[16] I. Y. Dodin and E. A. Startsev. "On applications of quantum computing to plasma simulations". *Physics of Plasmas* 28.9, 2021, p. 092101. DOI: 10.1063/5.0056974.

[17] W. Guan, G. Perdue, A. Pesah, M. Schuld, K. Terashi, S. Vallecorsa, and J.-R. Vlimant. "Quantum machine learning in high energy physics". *Machine Learning: Science and Technology* 2.1, 2021, p. 011003. DOI: 10.1088/2632-2153/abc17d.

[18] J. Preskill. "Simulating quantum field theory with a quantum computer". 2018. DOI: 10.48550/arXiv.1811.10085.

[19] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, et al. "Advances in quantum cryptography". *Advances in Optics and Photonics* 12.4, 2020, pp. 1012–1236. DOI: 10.1364/AOP.361502.

[20] M. Hillery, V. Bužek, and A. Berthiaume. "Quantum secret sharing". *Physical Review A* 59.3, 1999, p. 1829. DOI: 10.1103/PhysRevA.59.1829.

[21] Z.-J. Zhang. "Multiparty quantum secret sharing of secure direct communication". *Physics Letters A* 342.1-2, 2005, pp. 60–66. DOI: 10.1016/j.physleta.2005.05.049.

[22] A. Broadbent and C. Schaffner. "Quantum cryptography beyond quantum key distribution". *Designs, Codes and Cryptography* 78, 2016, pp. 351–382. DOI: 10.1007/s10623-015-0157-4.

[23] S. Wehner, D. Elkouss, and R. Hanson. "Quantum internet: A vision for the road ahead". *Science* 362.6412, 2018, eaam9288. DOI: 10.1126/science.aam9288.

[24] D. Gottesman, T. Jennewein, and S. Croke. "Longer-baseline telescopes using quantum repeaters". *Physical Review Letters* 109.7, 2012, p. 070503. DOI: 10.1103/PhysRevLett.109.070503.

[25] A. S. Cacciapuoti, M. Caleffi, F. Tafuri, F. S. Cataliotti, S. Gherardini, and G. Bianchi. "Quantum internet: Networking challenges in distributed quantum computing". *IEEE Network* 34.1, 2019, pp. 137–143. DOI: 10.1109/MNET.001.1900092.

[26] N. C. Jones, R. Van Meter, A. G. Fowler, P. L. McMahon, J. Kim, T. D. Ladd, and Y. Yamamoto. "Layered architecture for quantum computing". *Phys. Rev. X* 2, 3 2012, p. 031007. DOI: 10.1103/PhysRevX.2.031007.

[27] X. Fu, L. Riesebos, L. Lao, C. G. Almudever, F. Sebastiano, R. Versluis, E. Charbon, and K. Bertels. "A heterogeneous quantum computer architecture". *Proceedings of the ACM International Conference on Computing Frontiers*. 2016, pp. 323–330. DOI: 10.1145/2903150.2906827.

[28] J. M. Hornibrook et al. "Cryogenic control architecture for large-scale quantum computing". *Phys. Rev. Appl.* 3, 2 2015, p. 024010. DOI: 10.1103/PhysRevApplied.3.024010.

[29] A. Parra-Rodriguez, P. Lougovski, L. Lamata, E. Solano, and M. Sanz. "Digital-analog quantum computation". *Phys. Rev. A* 101, 2 2020, p. 022305. DOI: 10.1103/PhysRevA.101.022305.

[30] P. Kairys and T. S. Humble. "Parametrized Hamiltonian simulation using quantum optimal control". *Physical Review A* 104.4, 2021, p. 042602. DOI: 10.1103/PhysRevA.104.042602.

[31] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller. "Practical quantum advantage in quantum simulation". *Nature* 607.7920, 2022, pp. 667–676. DOI: 10.1038/s41586-022-04940-6.

[32] M. Mariantoni, H. Wang, T. Yamamoto, M. Neeley, R. C. Bialczak, Y. Chen, M. Lenander, E. Lucero, A. D. O'Connell, D. Sank, et al. "Implementing the quantum von Neumann architecture with superconducting circuits". *Science* 334.6052, 2011, pp. 61–65. DOI: 10.1126/science.1208517.

[33] P. Murali, N. M. Linke, M. Martonosi, A. J. Abhari, N. H. Nguyen, and C. H. Alderete. "Full-stack, real-system quantum computer studies: Architectural comparisons and design insights". *Proceedings of the 46th International Symposium on Computer Architecture.* 2019, pp. 527–540. DOI: 10.1145/3307650.3322273.

[34] K. Bharti et al. "Noisy intermediate-scale quantum algorithms". *Rev. Mod. Phys.* 94, 1 2022, p. 015004. DOI: 10.1103/RevModPhys.94.015004.

[35] K. A. Britt and T. S. Humble. "High-performance computing with quantum processing units". *ACM Journal on Emerging Technologies in Computing Systems* 13.3, 2017, pp. 1–13. DOI: 10.1145/3007651.

[36] A. McCaskey, E. Dumitrescu, D. Liakh, M. Chen, W. Feng, and T. Humble. "A language and hardware independent approach to quantum–classical computing". *SoftwareX* 7, 2018, pp. 245–254. DOI: 10.1016/j.softx.2018.07.007.

[37] T. S. Humble, A. McCaskey, D. I. Lyakh, M. Gowrishankar, A. Frisch, and T. Monz. "Quantum computers for high-performance computing". *IEEE Micro* 41.5, 2021, pp. 15–23. DOI: 10.1109/MM.2021.3099140.

[38] A. Dahlberg, B. van der Vecht, C. Delle Donne, M. Skrzypczyk, I. te Raa, W. Kozlowski, and S. Wehner. "NetQASM–A low-level instruction set architecture for hybrid quantum–classical programs in a quantum internet". *Quantum Science and Technology* 7.3, 2022, p. 035023. DOI: 10.1088/2058-9565/ac753f.

[39] R. Van Meter. *Quantum Networking.* ISTE. Wiley, 2014. DOI: 10.1002/9781118648919.

[40] D. Cuomo, M. Caleffi, and A. S. Cacciapuoti. "Towards a distributed quantum computing ecosystem". *IET Quantum Communication* 1.1, 2020, pp. 3–8. DOI: 10.1049/iet-qtc.2020.0002.

[41] M. Alshowkan, B. P. Williams, P. G. Evans, N. S. Rao, E. M. Simmerman, H.-H. Lu, N. B. Lingaraju, A. M. Weiner, C. E. Marvinney, Y.-Y. Pai, et al. "Reconfigurable quantum local area network over deployed fiber". *PRX Quantum* 2.4, 2021, p. 040304. DOI: 10.1103/PRXQuantum.2.040304.

[42] M. Alshowkan, P. G. Evans, B. P. Williams, N. S. Rao, C. E. Marvinney, Y.-Y. Pai, B. J. Lawrie, N. A. Peters, and J. M. Lukens. "Advanced architectures for high-performance quantum networking". *Journal of Optical Communications and Networking* 14.6, 2022, pp. 493–499. DOI: 10.1364/JOCN.450201.

[43] J. Chung, E. M. Eastman, G. S. Kanter, K. Kapoor, N. Lauk, C. H. Pena, R. K. Plunkett, N. Sinclair, J. M. Thomas, R. Valivarthi, et al. "Design and implementation of the Illinois Express quantum metropolitan area network". *IEEE Transactions on Quantum Engineering* 3, 2022, pp. 1–20. DOI: 10.1109/TQE.2022.3221029.

[44] D. Awschalom, K. K. Berggren, H. Bernien, S. Bhave, L. D. Carr, P. Davids, S. E. Economou, D. Englund, A. Faraon, M. Fejer, et al. "Development of quantum interconnects (QuICs) for next-generation information technologies". *PRX Quantum* 2.1, 2021, p. 017002. DOI: 10.1103/PRXQuantum.2.017002.

[45] V. Semenenko, X. Hu, E. Figueroa, and V. Perebeinos. "Entanglement generation in a quantum network with finite quantum memory lifetime". *AVS Quantum Science* 4.1, 2022, p. 012002. DOI: 10.1116/5.0082239.

[46] M. Alshowkan, J. M. Lukens, H.-H. Lu, B. T. Kirby, B. P. Williams, W. P. Grice, and N. A. Peters. "Broadband polarization-entangled source for C+L-band flex-grid quantum networks". *Optics Letters* 47.24, 2022, pp. 6480–6483. DOI: 10.1364/OL.471363.

[47] K. Azuma, S. Bäuml, T. Coopmans, D. Elkouss, and B. Li. "Tools for quantum network design". *AVS Quantum Science* 3.1, 2021, p. 014101. DOI: 10.1116/5.0024062.

[48] P. G. Evans, M. Alshowkan, D. Earl, D. D. Mulkey, R. Newell, G. Peterson, C. Safi, J. L. Tripp, and N. A. Peters. "Trusted node QKD at an electrical utility". *IEEE Access* 9, 2021, pp. 105220–105229. DOI: 10.1109/ACCESS.2021.3070222.

[49] M. Alshowkan, P. G. Evans, M. Starke, D. Earl, and N. A. Peters. "Authentication of smart grid communications using quantum key distribution". *Scientific Reports* 12.1, 2022, p. 12731. DOI: 10.1038/s41598-022-16090-w.

[50] B. A. Wilson, A. Miloshevsky, D. A. Hooper, W. Grice, and N. A. Peters. "Optical noise in a free-space quantum communications link from natural and nuclear disturbed environments". *New Journal of Physics* 24.6, 2022, p. 063035. DOI: 10.1088/1367-2630/ac77f4.

[51] J. Miguel-Ramiro, A. Pirker, and W. Dür. "Optimized quantum networks". *Quantum* 7, 2023, p. 919. DOI: 10.22331/q-2023-02-09-919.

[52] R. Valivarthi et al. "Teleportation systems toward a quantum internet". *PRX Quantum* 1, 2 2020, p. 020317. DOI: 10.1103/PRXQuantum.1.020317.

[53] J. Illiano, M. Caleffi, A. Manzalini, and A. S. Cacciapuoti. "Quantum internet protocol stack: A comprehensive survey". *Computer Networks* 213, 2022, p. 109092. DOI: 10.1016/j.comnet.2022.109092.

[54] J. Alnas, M. Alshowkan, N. S. Rao, N. A. Peters, and J. M. Lukens. "Optimal resource allocation for flexible-grid entanglement distribution networks". *Optics Express* 30.14, 2022, pp. 24375–24393. DOI: 10.1364/OE.458358.

[55] P. Shor. "Algorithms for quantum computation: Discrete logarithms and factoring". *Proceedings 35th Annual Symposium on Foundations of Computer Science.* 1994. DOI: 10.1109/sfcs.1994.365700.

[56] L. K. Grover. "A fast quantum mechanical algorithm for database search". *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing.* ACM Press, 1996. DOI: 10.1145/237814.237866.

[57] S. Jordan. "The quantum algorithm zoo". Available at https://quantumalgorithmzoo.org.

[58] S. Aaronson. "How much structure is needed for huge quantum speedups?" 2022. DOI: 10.48550/arXiv.2209.06930.

[59] A. M. Dalzell, A. W. Harrow, D. E. Koh, and R. L. La Placa. "How many qubits are needed for quantum computational supremacy?" *Quantum* 4, 2020, p. 264. DOI: 10.22331/q-2020-05-11-264.

[60] A. W. Harrow and A. Montanaro. "Quantum computational supremacy". *Nature* 549.7671, 2017, pp. 203–209. DOI: 10.1038/nature23458.

[61] J. Preskill. "Quantum computing and the entanglement frontier". 2012. DOI: 10.48550/arXiv.1203.5813.

[62] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, et al. "Variational quantum algorithms". *Nature Reviews Physics* 3.9, 2021, pp. 625–644. DOI: 10.1038/s42254-021-00348-9.

[63] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang. "Obstacles to variational quantum optimization from symmetry protection". *Phys. Rev. Lett.* 125, 26 2020, p. 260505. DOI: 10.1103/PhysRevLett.125.260505.

[64] C.-N. Chou, P. J. Love, J. S. Sandhu, and J. Shi. "Limitations of local quantum algorithms on random Max-k-XOR and beyond". *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).* Schloss Dagstuhl-Leibniz-Zentrum für Informatik. 2022. DOI: 10.4230/LIPIcs.ICALP.2022.41.

[65] E. Tang. "A quantum-inspired classical algorithm for recommendation systems". *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing.* 2019, pp. 217–228. DOI: 10.1145/3313276.3316310.

[66] N.-H. Chia, A. Gilyén, T. Li, H.-H. Lin, E. Tang, and C. Wang. "Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning". *Proceedings of the 52nd Annual ACM SIGACT symposium on theory of computing.* 2020, pp. 387–400. DOI: 10.1145/3357713.3384314.

[67] J. Cotler, H.-Y. Huang, and J. R. McClean. "Revisiting dequantization and quantum advantage in learning tasks". 2021. DOI: 10.48550/arXiv.2112.00811.

[68] S. Aaronson. "Read the fine print". *Nature Physics* 11.4, 2015, pp. 291–293. DOI: 10.1038/nphys3272.

[69] D. P. Williamson and D. B. Shmoys. *The design of approximation algorithms.* Cambridge University Press, 2011. DOI: 10.1017/CBO9780511921735.

[70] A. E. Feldmann, E. Lee, and P. Manurangsi. "A survey on approximation in parameterized complexity: Hardness and algorithms". *Algorithms* 13.6, 2020, p. 146. DOI: 10.3390/a13060146.

[71] T. Dokeroglu, E. Sevinc, T. Kucukyilmaz, and A. Cosar. "A survey on new generation metaheuristic algorithms". *Computers & Industrial Engineering* 137, 2019, p. 106040. DOI: https://doi.org/10.1016/j.cie.2019.106040.

[72] S. Gharibian, Y. Huang, Z. Landau, and S. W. Shin. "Quantum Hamiltonian complexity". *Foundations and Trends in Theoretical Computer Science* 10.3, 2015, pp. 159–282. DOI: 10.1561/0400000066.

[73] E. Lee. "Optimizing quantum circuit parameters via SDP". 2022. DOI: 10.48550/arXiv.2209.00789.

[74] R. King. "An improved approximation algorithm for quantum Max-Cut". 2022. DOI: 10.48550/arXiv.2209.02589.

[75] O. Parekh and K. Thompson. "An optimal product-state approximation for 2-local quantum Hamiltonians with positive terms". 2022. DOI: 10.48550/arXiv.2206.08342.

[76] A. Anshu, D. Gosset, K. J. M. Korol, and M. Soleimanifar. "Improved approximation algorithms for bounded-degree local Hamiltonians". *Physical Review Letters* 127.25, 2021, p. 250502. DOI: 10.1103/PhysRevLett.127.250502.

[77] Y. Hwang, J. Neeman, O. Parekh, K. Thompson, and J. Wright. "Unique games hardness of Quantum Max-Cut, and a conjectured vector-valued Borell's inequality". 2021. DOI: 10.48550/arXiv.2111.01254.

[78] F. G. Brandao and A. W. Harrow. "Product-state approximations to quantum states". *Communications in Mathematical Physics* 342, 2016, pp. 47–80. DOI: 10.1007/s00220-016-2575-1.

[79] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete. "Matrix product states and projected entangled pair states: concepts, symmetries, theorems". *Rev. Mod. Phys.* 93, 4 2021, p. 045003. DOI: 10.1103/RevModPhys.93.045003.

[80] C. Huang, F. Zhang, M. Newman, X. Ni, D. Ding, J. Cai, X. Gao, T. Wang, F. Wu, G. Zhang, et al. "Efficient parallelization of tensor network contraction for simulating quantum computation". *Nature Computational Science* 1.9, 2021, pp. 578–587. DOI: 10.1038/s43588-021-00119-7.

[81] T. Nguyen, D. Lyakh, E. Dumitrescu, D. Clark, J. Larkin, and A. McCaskey. "Tensor network quantum virtual machine for simulating quantum circuits at exascale". *ACM Transactions on Quantum Computing* 4.1, 2022. DOI: 10.1145/3547334.

[82] G. H. Booth, A. J. Thom, and A. Alavi. "Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space". *The Journal of Chemical Physics* 131.5, 2009, p. 054106. DOI: 10.1063/1.3193710.

[83] S. Zhang, J. Carlson, and J. E. Gubernatis. "Constrained path Monte Carlo method for fermion ground states". *Phys. Rev. B* 55, 12 1997, pp. 7464–7477. DOI: 10.1103/PhysRevB.55.7464.

[84] D. M. Ceperley and B. J. Alder. "Ground state of the electron gas by a stochastic method". *Phys. Rev. Lett.* 45, 7 1980, pp. 566–569. DOI: 10.1103/PhysRevLett.45.566.

[85] B. Heim, M. Soeken, S. Marshall, C. Granade, M. Roetteler, A. Geller, M. Troyer, and K. Svore. "Quantum programming languages". *Nature Reviews Physics* 2.12, 2020, pp. 709–722. DOI: 10.1038/s42254-020-00245-7.

[86] F. T. Chong, D. Franklin, and M. Martonosi. "Programming languages and compiler design for realistic quantum hardware". *Nature* 549.7671, 2017, pp. 180–187. DOI: 10.1038/nature23459.

[87] A. McCaskey, E. Dumitrescu, D. Liakh, and T. Humble. "Hybrid programming for near-term quantum computing systems". *International Conference on Rebooting Computing*. IEEE. 2018, pp. 1–12. DOI: 10.1109/ICRC.2018.8638598.

[88] D. S. Steiger, T. Häner, and M. Troyer. "ProjectQ: An open source software framework for quantum computing". *Quantum* 2, 2018, p. 49. DOI: 10.22331/q-2018-01-31-49.

[89] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan. "t|ket⟩: A retargetable compiler for NISQ devices". *Quantum Science and Technology* 6.1, 2020, p. 014003. DOI: 10.1088/2058-9565/ab8e92.

[90] A. McCaskey and T. Nguyen. "A MLIR dialect for quantum assembly languages". *International Conference on Quantum Computing and Engineering*. IEEE. 2021, pp. 255–264. DOI: 10.1109/QCE52317.2021.00043.

[91] A. J. McCaskey, D. I. Lyakh, E. F. Dumitrescu, S. S. Powers, and T. S. Humble. "XACC: A system-level software infrastructure for heterogeneous quantum–classical computing". *Quantum Science and Technology* 5.2, 2020, p. 024002. DOI: 10.1088/2058-9565/ab6bf6.

[92] A. McCaskey, T. Nguyen, A. Santana, D. Claudino, T. Kharazi, and H. Finkel. "Extending C++ for heterogeneous quantum-classical computing". *ACM Transactions on Quantum Computing* 2.2, 2021, pp. 1–36. DOI: 10.1145/3462670.

[93] A. JavadiAbhari, S. Patil, D. Kudrow, J. Heckey, A. Lvov, F. T. Chong, and M. Martonosi. "ScaffCC: Scalable compilation and analysis of quantum programs". *Parallel Computing* 45, 2015. Computing Frontiers 2014: Best Papers, pp. 2–17. DOI: 10.1016/j.parco.2014.12.001.

[94] A. Cross, A. Javadi-Abhari, T. Alexander, N. De Beaudrap, L. S. Bishop, S. Heidel, C. A. Ryan, P. Sivarajah, J. Smolin, J. M. Gambetta, et al. "OpenQASM 3: A broader and deeper quantum assembly language". *ACM Transactions on Quantum Computing* 3.3, 2022, pp. 1–50. DOI: 10.1145/3505636.

[95] D. Ittah, T. Häner, V. Kliuchnikov, and T. Hoefler. "QIRO: A static single assignment-based quantum program representation for optimization". *ACM Transactions on Quantum Computing* 3.3, 2022. DOI: 10.1145/3491247.

[96] R. S. Smith, M. J. Curtis, and W. J. Zeng. "A practical quantum instruction set architecture". 2016. DOI: 10.48550/arXiv.1608.03355.

[97] N. Killoran, J. Izaac, N. Quesada, V. Bergholm, M. Amy, and C. Weedbrook. "Strawberry Fields: A software platform for photonic quantum computing". *Quantum* 3, 2019, p. 129. DOI: 10.22331/q-2019-03-11-129.

[98] K. Singhal, K. Hietala, S. Marshall, and R. Rand. "Q# as a quantum algorithmic language". 2022. DOI: 10.48550/arXiv.2206.03532.

[99] J. R. McClean et al. "OpenFermion: The electronic structure package for quantum computers". *Quantum Science and Technology* 5.3, 2020, p. 034014. DOI: 10.1088/2058-9565/ab8ebc.

[100] T. Nguyen, L. Bassman Oftelie, P. C. Lotshaw, D. Lyakh, A. McCaskey, V. Leyton-Ortega, R. Pooser, W. Elwasif, T. S. Humble, and W. A. de Jong. "QuaSiMo: A composable library to program hybrid workflows for quantum simulation". *IET Quantum Communication* 2.4, 2021, pp. 160–170. DOI: 10.1049/qtc2.12024.

[101] W. Jang, K. Terashi, M. Saito, C. W. Bauer, B. Nachman, Y. Iiyama, R. Okubo, and R. Sawada. "Initial-state dependent optimization of controlled gate operations with quantum computer". *Quantum* 6, 2022, p. 798. DOI: 10.22331/q-2022-09-08-798.

[102] Y. Shi et al. "Resource-efficient quantum computing by breaking abstractions". *Proceedings of the IEEE* 108.8, 2020, pp. 1353–1370. DOI: 10.1109/JPROC.2020.2994765.

[103] E. Wilson, S. Singh, and F. Mueller. "Just-in-time quantum circuit transpilation reduces noise". *IEEE International Conference on Quantum Computing and Engineering*. 2020, pp. 345–355. DOI: 10.1109/QCE49297.2020.00050.

[104] E. Smith, M. G. Davis, J. Larson, E. Younis, L. B. Oftelie, W. Lavrijsen, and C. Iancu. "LEAP: Scaling numerical optimization based synthesis using an incremental approach". *ACM Transactions on Quantum Computing* 4.1, 2023, pp. 1–23. DOI: 10.1145/3548693.

[105] E. Younis and C. Iancu. "Quantum circuit optimization and transpilation via parameterized circuit instantiation". *IEEE International Conference on Quantum Computing and Engineering*. IEEE Computer Society, 2022, pp. 465–475. DOI: 10.1109/QCE53715.2022.00068.

[106] T. Patel, E. Younis, C. Iancu, W. de Jong, and D. Tiwari. "QUEST: Systematically approximating quantum circuits for higher output fidelity". *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS '22. Association for Computing Machinery, 2022, 514—528. DOI: 10.1145/3503222.3507739.

[107] V. Kliuchnikov, D. Maslov, and M. Mosca. "Fast and efficient exact synthesis of single-qubit unitaries generated by Clifford and T gates". *Quantum Information and Computation* 13.7–8, 2013, pp. 607–630. DOI: 10.26421/qic13.7-8-4.

[108] M. Amy, D. Maslov, M. Mosca, and M. Roetteler. "A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits". *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 32.6, 2013, pp. 818–830. DOI: 10.1109/TCAD.2013.2244643.

[109] B. Giles and P. Selinger. "Exact synthesis of multiqubit Clifford+T circuits". *Phys. Rev. A* 87, 3 2013, p. 032332. DOI: 10.1103/PhysRevA.87.032332.

[110] R. Wille, S. Hillmich, and L. Burgholzer. "Efficient and correct compilation of quantum circuits". *IEEE International Symposium on Circuits and Systems*. 2020, pp. 1–5. DOI: 10.1109/iscas45731.2020.9180791.

[111] R. R. Tucci. "Qubiter algorithm modification, expressing unstructured unitary matrices with fewer CNOTs". 2004. DOI: 10.48550/arXiv.quant-ph/0411027.

[112] S. Pakin. "A quantum macro assembler". *IEEE High Performance Extreme Computing Conference*. 2016, pp. 1–8. DOI: 10.1109/HPEC.2016.7761637.

[113] B. Tan and J. Cong. "Optimal layout synthesis for quantum computing". *Proceedings of the 39th International Conference on Computer-Aided Design*. ACM, 2020. DOI: 10.1145/3400302.3415620.

[114] G. Li, Y. Ding, and Y. Xie. "Tackling the qubit mapping problem for NISQ-era quantum devices". *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2019. DOI: 10.1145/3297858.3304023.

[115] S. S. Tannu and M. K. Qureshi. "Not all qubits are created equal". *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2019. DOI: 10.1145/3297858.3304007.

[116] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi. "Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers". *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2019. DOI: 10.1145/3297858.3304075.

[117] G. Nannicini, L. S. Bishop, O. Günlük, and P. Jurcevic. "Optimal qubit assignment and routing via integer programming". *ACM Transactions on Quantum Computing* 4.1, 2022. DOI: 10.1145/3544563.

[118] P. Gokhale, A. Javadi-Abhari, N. Earnest, Y. Shi, and F. T. Chong. "Optimized quantum compilation for near-term algorithms with OpenPulse". *53rd Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, 2020, pp. 186–200. DOI: 10.1109/micro50266.2020.00027.

[119] H. Ball et al. "Software tools for quantum control: Improving quantum computer performance through noise and error suppression". *Quantum Science and Technology* 6.4, 2021, p. 044011. DOI: 10.1088/2058-9565/abdca6.

[120] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, et al. "Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe". *EPJ Quantum Technology* 9.1, 2022, p. 19. DOI: 10.1140/epjqt/s40507-022-00138-x.

[121] M. A. Perlin, Z. H. Saleem, M. Suchara, and J. C. Osborn. "Quantum circuit cutting with maximum-likelihood tomography". *npj Quantum Information* 7.1, 2021. DOI: 10.1038/s41534-021-00390-6.

[122] W. Tang, T. Tomesh, M. Suchara, J. Larson, and M. Martonosi. "CutQC: Using small quantum computers for large quantum circuit evaluations". *Proceedings of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 2021, pp. 473–486. DOI: 10.1145/3445814.3446758.

[123] A. S. Nikolaeva, E. O. Kiktenko, and A. K. Fedorov. "Compiling quantum circuits with qubits embedded in trapped-ion quidts". 2023. DOI: 10.48550/arXiv.2302.02966.

[124] T. Häner, D. S. Steiger, T. Hoefler, and M. Troyer. "Distributed quantum computing with QMPI". *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. SC '21. Association for Computing Machinery, 2021. DOI: 10.1145/3458817.3476172.

[125] N. Goss et al. "High-fidelity qutrit entangling gates for superconducting circuits". *Nature Communications* 13.1, 2022, p. 7481. DOI: 10.1038/s41467-022-34851-z.

[126] P. Gokhale, J. M. Baker, C. Duckering, N. C. Brown, K. R. Brown, and F. T. Chong. "Asymptotic improvements to quantum circuits via qutrits". *Proceedings of the 46th International Symposium on Computer Architecture*. ISCA '19. Association for Computing Machinery, 2019, 554–566. DOI: 10.1145/3307650.3322253.

[127] R. LaRose, A. Mari, S. Kaiser, P. J. Karalekas, A. A. Alves, P. Czarnik, M. El Mandouh, M. H. Gordon, Y. Hindy, A. Robertson, et al. "Mitiq: A software package for error mitigation on noisy quantum computers". *Quantum* 6, 2022, p. 774. DOI: 10.22331/q-2022-08-11-774.

[128] A. J. Landahl and B. C. A. Morrison. "Logical Majorana fermions for fault-tolerant quantum simulation". 2021. DOI: 10.48550/arXiv.2110.10280.

[129] Google Quantum AI. "Suppressing quantum errors by scaling a surface code logical qubit". *Nature* 614.7949, 2023, p. 676. DOI: 10.1038/s41586-022-05434-1.

[130] A. Holmes, M. R. Jokar, G. Pasandi, Y. Ding, M. Pedram, and F. T. Chong. "NISQ+: Boosting quantum computing power by approximating quantum error correction". 2020. DOI: 10.48550/arXiv.2004.04794.

[131] A. Jayashankar and P. Mandayam. "Quantum error correction: Noise-adapted techniques and applications". 2022. DOI: 10.48550/arXiv.2208.00365.

[132] S. Bravyi, M. Englbrecht, R. König, and N. Peard. "Correcting coherent errors with surface codes". *npj Quantum Information* 4.1, 2018, p. 55. DOI: 10.1038/s41534-018-0106-y.

[133] T. Grurl, C. Pichler, J. Fuß, and R. Wille. "Automatic implementation and evaluation of error-correcting codes for quantum computing: An open-source framework for quantum error correction". 2023. DOI: 10.48550/arXiv.2301.05731.

[134] D. P. DiVincenzo. "The physical implementation of quantum computation". *Fortschritte der Physik: Progress of Physics* 48.9-11, 2000, pp. 771–783.

[135] D. A. Lidar and T. A. Brun, eds. *Quantum Error Correction*. Cambridge University Press, 2013. DOI: 10.1017/CBO9781139034807.

[136] D. A. Lidar. "Review of decoherence-free subspaces, noiseless subsystems, and dynamical decoupling". *Advances in Chemical Physics*. John Wiley & Sons, Inc., 2014, pp. 295–354. DOI: 10.1002/9781118742631.ch11.

[137] P. W. Shor. "Scheme for reducing decoherence in quantum computer memory". *Physical Review A* 52.4, 1995, R2493. DOI: 10.1103/PhysRevA.52.R2493.

[138] D. Gottesman. "Stabilizer codes and quantum error correction". 1997. DOI: 10.48550/arXiv.quant-ph/9705052.

[139] B. M. Terhal. "Quantum error correction for quantum memories". *Reviews of Modern Physics* 87.2, 2015, p. 307. DOI: 10.1103/RevModPhys.87.307.

[140] D. A. Lidar, I. L. Chuang, and K. B. Whaley. "Decoherence-free subspaces for quantum computation". *Physical Review Letters* 81.12, 1998, p. 2594. DOI: 10.1103/PhysRevLett.81.2594.

[141] J. Kempe, D. Bacon, D. A. Lidar, and K. B. Whaley. "Theory of decoherence-free fault-tolerant universal quantum computation". *Phys. Rev. A* 63, 4 2001, p. 042307. DOI: 10.1103/PhysRevA.63.042307.

[142] L. Viola, E. Knill, and S. Lloyd. "Dynamical decoupling of open quantum systems". *Physical Review Letters* 82.12, 1999, p. 2417. DOI: 10.1103/PhysRevLett.82.2417.

[143] S. Endo, S. C. Benjamin, and Y. Li. "Practical quantum error mitigation for near-future applications". *Physical Review X* 8.3, 2018. DOI: 10.1103/physrevx.8.031027.

[144] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta. "Error mitigation extends the computational reach of a noisy quantum processor". *Nature* 567.7749, 2019, pp. 491–495. DOI: 10.1038/s41586-019-1040-7.

[145] D. M. Debroy and K. R. Brown. "Extended flag gadgets for low-overhead circuit verification". *Phys. Rev. A* 102, 5 2020, p. 052409. DOI: 10.1103/PhysRevA.102.052409.

[146] A. Gonzales, R. Shaydulin, Z. H. Saleem, and M. Suchara. "Quantum error mitigation by Pauli check sandwiching". *Scientific Reports* 13.1, 2023, p. 2122. DOI: 10.1038/s41598-023-28109-x.

[147] C. Piveteau and D. Sutter. "Circuit knitting with classical communication". 2022. DOI: 10.48550/arxiv.2205.00016.

[148] J. Liu, A. Gonzales, and Z. H. Saleem. "Classical simulators as quantum error mitigators via circuit cutting". 2022. DOI: 10.48550/arXiv.2212.07335.

[149] C. Ospelkaus, C. E. Langer, J. M. Amini, K. R. Brown, D. Leibfried, and D. J. Wineland. "Trapped-ion quantum logic gates based on oscillating magnetic fields". *Phys. Rev. Lett.* 101.9, 2008, p. 090502. DOI: 10.1103/PhysRevLett.101.090502.

[150] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver. "Noise spectroscopy through dynamical decoupling with a superconducting flux qubit". *Nature Physics* 7.7, 2011, pp. 565–570. DOI: 10.1038/nphys1994.

[151] V. Frey, S. Mavadia, L. Norris, W De Ferranti, D Lucarelli, L Viola, and M. Biercuk. "Application of optimal band-limited control protocols to quantum noise sensing". *Nature Communications* 8.1, 2017, pp. 1–8. DOI: 10.1038/s41467-017-02298-2.

[152] D. Buterakos, R. E. Throckmorton, and S. D. Sarma. "Crosstalk error correction through dynamical decoupling of single-qubit gates in capacitively coupled singlet-triplet semiconductor spin qubits". *Physical Review B* 97.4, 2018. DOI: 10.1103/physrevb.97.045431.

[153] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. "A quantum engineer's guide to superconducting qubits". *Applied Physics Reviews* 6.2, 2019, p. 021318. DOI: 10.1063/1.5089550.

[154] A. Ash-Saki, M. Alam, and S. Ghosh. "Experimental characterization, modeling, and analysis of crosstalk in a quantum computer". *IEEE Transactions on Quantum Engineering* 1, 2020, pp. 1–6. DOI: 10.1109/TQE.2020.3023338.

[155] J. Ku, X. Xu, M. Brink, D. C. McKay, J. B. Hertzberg, M. H. Ansari, and B. L. T. Plourde. "Suppression of unwanted ZZ interactions in a hybrid two-qubit system". *Phys. Rev. Lett.* 125, 20 2020, p. 200504. DOI: 10.1103/PhysRevLett.125.200504.

[156] W. I. L. Lawrie, M. Russ, F. van Riggelen, N. W. Hendrickx, S. L. de Snoo, A. Sammak, G. Scappucci, and M. Veldhorst. "Simultaneous driving of semiconductor spin qubits at the fault-tolerant threshold". 2021. DOI: 10.48550/arxiv.2109.07837.

[157] V. Frey, L. M. Norris, L. Viola, and M. J. Biercuk. "Simultaneous spectral estimation of dephasing and amplitude noise on a qubit sensor via optimally band-limited control". *Phys. Rev. Applied* 14, 2 2020, p. 024021. DOI: 10.1103/PhysRevApplied.14.024021.

[158] J. He, Q. Liu, Z. Yang, Q. Niu, X. Ban, and J. Wang. "Noise spectroscopy with a Rydberg ensemble in a hot atomic vapor cell". *Phys. Rev. A* 104, 6 2021, p. 063120. DOI: 10.1103/PhysRevA.104.063120.

[159] P. Parrado-Rodríguez, C. Ryan-Anderson, A. Bermudez, and M. Müller. "Crosstalk suppression for fault-tolerant quantum error correction with trapped ions". *Quantum* 5, 2021, p. 487. DOI: 10.22331/q-2021-06-29-487.

[160] E. J. Connors, J Nelson, L. F. Edge, and J. M. Nichol. "Charge-noise spectroscopy of Si/SiGe quantum dots via dynamically-decoupled exchange oscillations". *Nature Communications* 13.1, 2022, pp. 1–9. DOI: 10.1038/s41467-022-28519-x.

[161] I. Heinz and G. Burkard. "Crosstalk analysis for simultaneously driven two-qubit gates in spin qubit arrays". *Phys. Rev. B* 105, 8 2022, p. 085414. DOI: 10.1103/PhysRevB.105.085414.

[162] C. Fang, Y. Wang, S. Huang, K. R. Brown, and J. Kim. "Crosstalk suppression in individually addressed two-qubit gates in a trapped-ion quantum computer". *Phys. Rev. Lett.* 129, 24 2022, p. 240504. DOI: 10.1103/PhysRevLett.129.240504.

[163] D. W. Leung, M. A. Nielsen, I. L. Chuang, and Y. Yamamoto. "Approximate quantum error correction can lead to better codes". *Phys. Rev. A* 56, 4 1997, pp. 2567–2573. DOI: 10.1103/PhysRevA.56.2567.

[164] B. Schumacher and M. D. Westmoreland. "Approximate quantum error correction". *Quantum Information Processing* 1, 2002, pp. 5–12. DOI: 10.1023/A:1019653202562.

[165] C. Bény and O. Oreshkov. "General conditions for approximate quantum error correction and near-optimal recovery channels". *Phys. Rev. Lett.* 104, 12 2010, p. 120501. DOI: 10.1103/PhysRevLett.104.120501.

[166] D. M. Debroy, L. Egan, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Cetina, C. Monroe, and K. R. Brown. "Optimizing stabilizer parities for improved logical qubit memories". *Phys. Rev. Lett.* 127, 24 2021, p. 240501. DOI: 10.1103/PhysRevLett.127.240501.

[167] D. Greenbaum. "Introduction to quantum gate set tomography". 2015. DOI: 10.48550/arXiv.1509.02921.

[168] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout. "Gate set tomography". *Quantum* 5, 2021, p. 557. DOI: 10.22331/q-2021-10-05-557.

[169] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. "Randomized benchmarking of quantum gates". *Phys. Rev. A* 77, 1 2008, p. 012307. DOI: 10.1103/PhysRevA.77.012307.

[170] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt. "Characterizing large-scale quantum computers via cycle benchmarking". *Nature Communications* 10.1, 2019, p. 5347. DOI: 10.1038/s41467-019-13068-7.

[171] G. A. Paz-Silva, L. M. Norris, and L. Viola. "Multiqubit spectroscopy of gaussian quantum noise". *Phys. Rev. A* 95, 2 2017, p. 022121. DOI: 10.1103/PhysRevA.95.022121.

[172] K. Khodjasteh and L. Viola. "Dynamically error-corrected gates for universal quantum computation". *Phys. Rev. Lett.* 102, 8 2009, p. 080501. DOI: 10.1103/PhysRevLett.102.080501.

[173] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger. "Optimized dynamical decoupling in a model quantum memory". *Nature* 458.7241, 2009, pp. 996–1000. DOI: 10.1038/nature07951.

[174] H. Ball and M. J. Biercuk. "Walsh-synthesized noise filters for quantum logic". *EPJ Quantum Technology* 2.1, 2015, pp. 1–45. DOI: 10.1140/epjqt/s40507-015-0022-4.

[175] Y. Oda, D. Lucarelli, K. Schultz, B. D. Clader, and G. Quiroz. "Optimally band-limited noise filtering for single-qubit gates". *Phys. Rev. Appl.* 19, 1 2023, p. 014062. DOI: 10.1103/PhysRevApplied.19.014062.

[176] L.-A. Wu and D. A. Lidar. "Creating decoherence-free subspaces using strong and fast pulses". *Phys. Rev. Lett.* 88, 20 2002, p. 207902. DOI: 10.1103/PhysRevLett.88.207902.

[177] A. Shabani and D. A. Lidar. "Theory of initialization-free decoherence-free subspaces and subsystems". *Phys. Rev. A* 72, 4 2005, p. 042303. DOI: 10.1103/PhysRevA.72.042303.

[178] L.-A. Wu and M. S. Byrd. “Self-protected quantum algorithms based on quantum state tomography”. *Quantum Inf Process* 8, 2009, pp. 1–12. DOI: 10.1007/s11128-008-0090-9.

[179] X. Wang, M. Byrd, and K. Jacobs. “Minimal noise subsystems”. *Phys. Rev. Lett.* 116, 9 2016, p. 090404. DOI: 10.1103/PhysRevLett.116.090404.

[180] A. Zlokapa and A. Gheorghiu. “A deep learning model for noise prediction on near-term quantum devices”. 2020. DOI: 10.48550/arXiv.2005.10811.

[181] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio. “Error mitigation with Clifford quantum-circuit data”. *Quantum* 5, 2021, p. 592. DOI: 10.22331/q-2021-11-26-592.

[182] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li. “Learning-based quantum error mitigation”. *PRX Quantum* 2.4, 2021. DOI: 10.1103/prxquantum.2.040330.

[183] K. Schultz, R. LaRose, A. Mari, G. Quiroz, N. Shammah, B. D. Clader, and W. J. Zeng. “Impact of time-correlated noise on zero-noise extrapolation”. *Physical Review A* 106.5, 2022, p. 052406. DOI: 10.1103/PhysRevA.106.052406.

[184] Google Quantum AI. “Exponential suppression of bit or phase errors with cyclic error correction”. *Nature* 595.7867, 2021, pp. 383–387. DOI: 10.1038/s41586-021-03588-y.

[185] Y. Zhao, Y. Ye, H.-L. Huang, Y. Zhang, D. Wu, H. Guan, Q. Zhu, Z. Wei, T. He, S. Cao, et al. “Realization of an error-correcting surface code with superconducting qubits”. *Physical Review Letters* 129.3, 2022, p. 030501. DOI: 10.1103/PhysRevLett.129.030501.

[186] T. Kurita, H. Qassim, M. Ishii, H. Oshima, S. Sato, and J. Emerson. “Synergetic quantum error mitigation by randomized compiling and zero-noise extrapolation for the variational quantum eigensolver”. 2022. DOI: 10.48550/arXiv.2212.11198.

[187] S. Guo et al. “Scalable quantum computational chemistry with superconducting qubits”. 2022. DOI: 10.48550/arXiv.2212.08006.

[188] Y. Kim, C. J. Wood, T. J. Yoder, S. T. Merkel, J. M. Gambetta, K. Temme, and A. Kandala. “Scalable error mitigation for noisy quantum circuits produces competitive expectation values”. *Nature Physics*, 2023, pp. 1–8. DOI: 10.1038/s41567-022-01914-3.

[189] A. Zaman and H. Y. Wong. “Study of error propagation and generation in Harrow–Hassidim–Lloyd (HHL) quantum algorithm”. *Latin American Electron Devices Conference*. IEEE. 2022, pp. 1–4. DOI: 10.1109/LAEDC54796.2022.9908231.

[190] Z. Yu and Y. Li. “Analysis of error propagation in quantum computers”. 2022. DOI: 10.48550/arXiv.2209.01699.

[191] G. González-García, R. Trivedi, and J. I. Cirac. “Error propagation in NISQ devices for solving classical optimization problems”. *PRX Quantum* 3, 4 2022, p. 040326. DOI: 10.1103/PRXQuantum.3.040326.

[192] S. Flannigan, N. Pearson, G. H. Low, A Buyskikh, I. Bloch, P. Zoller, M. Troyer, and A. J. Daley. “Propagation of errors and quantitative quantum simulation with quantum advantage”. *Quantum Science and Technology* 7.4, 2022, p. 045025. DOI: 10.1088/2058-9565/ac88f5.

[193] R. S. Bennink, E. M. Ferragut, T. S. Humble, J. A. Laska, J. J. Nutaro, M. G. Pleszkoch, and R. C. Pooser. “Unbiased simulation of near-Clifford quantum circuits”. *Physical Review A* 95.6, 2017, p. 062337. DOI: 10.1103/PhysRevA.95.062337.

[194] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard. “Simulation of quantum circuits by low-rank stabilizer decompositions”. *Quantum* 3, 2019, p. 181. DOI: 10.22331/q-2019-09-02-181.

[195] B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T. S. Humble, R. Biswas, E. G. Rieffel, A. Ho, and S. Mandrà. “Establishing the quantum supremacy frontier with a 281 Pflop/s simulation”. *Quantum Science and Technology* 5.3, 2020, p. 034003. DOI: 10.1088/2058-9565/ab7eeb.

[196] T. Jones, A. Brown, I. Bush, and S. C. Benjamin. “QuEST and high performance simulation of quantum computers”. *Scientific Reports* 9.1, 2019, pp. 1–11. DOI: 10.1038/s41598-019-47174-9.

[197] A. Li, B. Fang, C. Granade, G. Prawiroatmodjo, B. Heim, M. Roetteler, and S. Krishnamoorthy. “SV-sim: Scalable PGAS-based state vector simulation of quantum circuits”. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 2021, pp. 1–14. DOI: 10.1145/3458817.3476169.

[198] B. Li, S. Ahmed, S. Saraogi, N. Lambert, F. Nori, A. Pitchford, and N. Shammah. “Pulse-level noisy quantum circuits with QuTiP”. *Quantum* 6, 2022, p. 630. DOI: 10.22331/q-2022-01-24-630.

[199] T. Tomesh, P. Gokhale, V. Omole, G. S. Ravi, K. N. Smith, J. Viszlai, X.-C. Wu, N. Hardavellas, M. R. Martonosi, and F. T. Chong. "SupermarQ: A scalable quantum benchmark suite". *International Symposium on High-Performance Computer Architecture*. IEEE. 2022, pp. 587–603. DOI: 10.1109/HPCA53966.2022.00050.

[200] A. Li, S. Stein, S. Krishnamoorthy, and J. Ang. "QASMBench: A low-level quantum benchmark suite for NISQ evaluation and simulation". *ACM Transactions on Quantum Computing*, 2022. DOI: 10.1145/3550488.

[201] J. Golden, A. Bärtschi, D. O'Malley, and S. Eidenbenz. "Fair sampling error analysis on NISQ devices". *ACM Transactions on Quantum Computing* 3.2, 2022. DOI: 10.1145/3510857.

[202] M. Suchara, J. Kubiatowicz, A. Faruque, F. T. Chong, C.-Y. Lai, and G. Paz. "QuRE: The quantum resource estimator toolbox". *31st International Conference on Computer Design*. IEEE. 2013, pp. 419–426. DOI: 10.1109/iccd.2013.6657074.

[203] M. Grassl, B. Langenberg, M. Roetteler, and R. Steinwandt. "Applying Grover's algorithm to AES: Quantum resource estimates". *Post-Quantum Cryptography*. Springer. 2016, pp. 29–43. DOI: 10.1007/978-3-319-29360-8_3.

[204] M. Roetteler, M. Naehrig, K. M. Svore, and K. Lauter. "Quantum resource estimates for computing elliptic curve discrete logarithms". *Advances in Cryptology*. Springer. 2017, pp. 241–270. DOI: 10.1007/978-3-319-70697-9_9.

[205] G. H. Low, N. P. Bauman, C. E. Granade, B. Peng, N. Wiebe, E. J. Bylaska, D. Wecker, S. Krishnamoorthy, M. Roetteler, K. Kowalski, et al. "Q# and NWChem: Tools for scalable quantum chemistry on quantum computers". 2019. DOI: 10.48550/arXiv.1904.01131.

[206] O. Di Matteo, V. Gheorghiu, and M. Mosca. "Fault-tolerant resource estimation of quantum random-access memories". *IEEE Transactions on Quantum Engineering* 1, 2020, pp. 1–13. DOI: 10.1109/tqe.2020.2965803.

[207] J. F. Gonthier, M. D. Radin, C. Buda, E. J. Doskocil, C. M. Abuan, and J. Romero. "Measurements as a roadblock to near-term practical quantum advantage in chemistry: Resource analysis". *Physical Review Research* 4.3, 2022. DOI: 10.1103/physrevresearch.4.033154.

[208] Z. Cai. "Resource estimation for quantum variational simulations of the Hubbard model". *Phys. Rev. Appl.* 14, 1 2020, p. 014059. DOI: 10.1103/PhysRevApplied.14.014059.

[209] A. Paler and R. Basmadjian. "Energy cost of quantum circuit optimisation: Predicting that optimising Shor's algorithm circuit uses 1 GWh". *ACM Transactions on Quantum Computing* 3.1, 2022. DOI: 10.1145/3490172.

[210] J. Helsen, I. Roth, E. Onorati, A. Werner, and J. Eisert. "General framework for randomized benchmarking". *PRX Quantum* 3, 2 2022, p. 020357. DOI: 10.1103/PRXQuantum.3.020357.

[211] X. Fu et al. "Quingo: A programming framework for heterogeneous quantum-classical computing with NISQ features". *ACM Transactions on Quantum Computing* 2.4, 2021. DOI: 10.1145/3483528.

[212] S. Dasgupta and T. S. Humble. "Characterizing the reproducibility of noisy quantum circuits". *Entropy* 24.2, 2022, p. 244. DOI: 10.3390/e24020244.

[213] Microsoft. "Azure Quantum Resource Estimator". https://aka.ms/AQ/RE. 2022.

Quantum State Passing Interfaces for Distributed Quantum Computing

Bruno Abreu[1,*] – babreu@illinois.edu
Santiago Núñez-Corrales[1,2] – nunezco2@illinois.edu
[1]National Center for Supercomputing Applications
[2]Illinois Quantum Information Science and Technology Center
University of Illinois Urbana-Champaign
[*]*Corresponding author*

Topics: applications; models; codesign and integration;

Challenge: Distributed quantum computing constitutes a scalable path to achieving an overall number of logical qubits that enable real-world quantum computing applications to flourish [1]. Given increasing evidence on possible scaling limits of entanglement [2] and the many technical challenges involved in establishing fully quantum interconnects across multiple physically distant devices [3], we understand that practical solutions must consider using hybrid classical-quantum interconnects to support near-to-mid-term distributed quantum computing effectively [4]. At the same time, we believe that programmers and practitioners should not necessarily be directly exposed to such details when they already have the difficult task of learning an intricate programming model that relies on the yet poor understanding of basic quantum mechanics over the research computing community. A successful journey toward large-scale quantum computing applications must therefore encompass programming models and frameworks that abstract most of the subtleties of interconnected heterogeneous quantum devices away from end users, allowing them to distribute computation similarly to what has been done for decades in classical computing while keeping those details accessible to researchers focused on developing quantum networks.

Opportunity: Because quantum computing technologies are still in their early stages, programmers and application developers are often forced to work on a level that, in a sense, is exceedingly close to quantum microcode or below, expressed through the ubiquitous circuit- or gate-based programming model [5]. Naturally, broad adoption of quantum computing frameworks will require abstraction layers that allow them to shift their focus to the application level. In this context, distributed quantum computing is no different and presents the opportunity to create and disseminate programming models designed to encapsulate features of quantum machines and interconnected quantum systems in general. By collaborating and codesigning a standard and abstract framework for hybrid quantum-classical communication and distribution, the quantum computing community will provide insights, tools, and opportunities for researchers and practitioners to work on both the application development front and the hardware and architecture design one. This route has the potential to ultimately lead to faster development of practical quantum computing applications and pave the way for larger-scale quantum computing in the future. By focusing on the information needed to pass quantum states rather than qubits or gates, such a framework can be flexible enough to incorporate situations where already existing classical interconnects can be used to transmit protocols that regenerate local superposition and entanglement with, for example, trace stacks of operations performed by distant quantum processors, which can then easily facilitate applications where many-qubit entanglement is not strictly necessary. Concurrently, short-distance quantum networks can create entanglement across clustered quantum computing nodes, which are programmed using the same interface, setting up the stage for seamless hybrid distributed quantum computing. We also foresee distributed quantum computing extending the lifetime of NISQ and intermediate devices by effectively pooling their resources, thus increasing long-term investment in the field.

Assessment: The primary measure of success of this effort is the assemblage of a community or forum with the mission of building an abstract programming model and framework for hybrid communication-based distributed quantum computing, possibly resembling what the MPI Forum has established for classical computing. Potential solutions should be evaluated based on their flexibility, the potential for widespread adoption, and the capacity to simultaneously make implementation details transparent to application developers and opaque to hardware, architecture, and network researchers; in essence, we expect the largest impact of this technology to be a substantial increase in quantum programs capable of using distributed resources, as well as decrease code complexity. Resources that will support the development effort include efficient quantum simulators for prototyping, partitionable quantum computers with a substantial number of qubits (hundreds) for testing, and reliable (lossy) quantum transmission lines for short (large) distance interconnects between devices for final implementation. Artificial intelligence will likely be an essential resource in deciding which level of hybridization to use for communication.

Timeliness: From its inception to accomplishing the goal of enabling exascale computing, the MPI standard endured three decades of intense development. The quantum computing ecosystem evolves quickly. Tools, frameworks, and programming models are increasingly intricate, creating valleys of developing communities that hinder the communication of ideas across this landscape. Forging a community dedicated to establishing a *de facto* standard for distributed quantum computing is a timely effort that will ultimately lead to scalable solutions across multiple heterogeneous platforms, unleashing the power of quantum computing to real-world applications.

References

[1] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, "Practical quantum advantage in quantum simulation," *Nature*, vol. 607, pp. 667–676, 7 2022.

[2] C.-c. Liu, D. Wang, W.-y. Sun, and L. Ye, "Quantum fisher information, quantum entanglement and correlation close to quantum critical phenomena," *Quantum Information Processing*, vol. 16, pp. 1–15, 2017.

[3] D. D. Awschalom *et al.*, "A roadmap for quantum interconnects," *Report ANL-22/83 179439, OSTI 1900586, USDOE U.S. Department of Energy*, 7 2022.

[4] G.-L. Long, D. Pan, Y.-B. Sheng, Q. Xue, J. Lu, and L. Hanzo, "An evolutionary pathway for the quantum internet relying on secure classical repeaters," *IEEE Network*, vol. 36, pp. 82–88, 5 2022.

[5] M. A. Serrano, J. A. Cruz-Lemus, R. Perez-Castillo, and M. Piattini, "Quantum software components and platforms: Overview and quality assessment," *ACM Computing Surveys*, vol. 55, pp. 1–31, 8 2023.

Scalable quantum chemistry algorithms

PI: Yuri Alexeev, Principal Project Specialist, Argonne National Laboratory, yuri@anl.gov
Co-PI: Matthew Otten, Scientist, HRL Laboratories
Co-PI: Stephen Gray, Senior Scientist, Argonne National Laboratory
Co-PI: Laura Gagliardi, Professor, University of Chicago

Call: Position Paper for Workshop on Quantum Computing and Networking
Topic: Algorithms

Challenge: Using quantum chemistry to compute the properties of large chemical systems is important because these systems play essential roles in many physical, chemical and biological processes. For example, studying large biomolecules such as proteins and DNA can help in understanding their functions and interactions, which is critical for developing new drugs and therapies. Many other materials such as polymers, ceramics, and semiconductors have complex structures that are difficult to analyze experimentally. Simulating these materials using quantum chemistry can provide insight into their electronic and structural properties, enabling the design of new materials with improved properties.

Both classical and quantum computers face challenges when it comes to computing large chemical systems. The fundamental reason is that simulating quantum systems accurately on classical computers requires an exponential amount of effort, which quickly becomes impractical as the size of the system increases. Quantum computers can in principle perform these calculations more efficiently than classical computers, but current quantum hardware is limited in terms of the number of qubits and the coherence time of these qubits. This makes it difficult to simulate large chemical systems accurately on current quantum hardware. Here is a detailed list of current problems:

1. **Quantum computer hardware:** Currently available quantum computers have limited qubit counts and high error rates, making it difficult to perform accurate quantum chemistry calculations.
2. **Quantum software:** Developing quantum algorithms and software for quantum chemistry calculations is a complex task and there is a need for better software tools and frameworks to make it easier for researchers to develop and run quantum chemistry calculations on quantum computers.
3. **Mapping quantum chemistry algorithms to quantum hardware:** There is a need to map quantum chemistry algorithms to the hardware architecture of quantum computers, which is a challenging task due to differences in the hardware architectures of different quantum computing platforms.
4. **Limited qubit connectivity:** Quantum chemistry calculations often require qubit connectivity that is not available on current quantum computing hardware, leading to difficulties, such as increased gate depth, in implementing the necessary quantum circuits.
5. **Accuracy of quantum chemistry calculations:** Noise and errors in quantum hardware can have a significant impact on the accuracy of quantum chemistry calculations creating a need for better error correction and error mitigation techniques.
6. **Classical post-processing:** Even after performing quantum chemistry calculations on a quantum computer, classical post-processing is often required to obtain meaningful results. This can be a complex and computationally intensive task.
7. **Lack of standardization:** There is currently a lack of standardization in quantum chemistry calculations on quantum computers, making it difficult to compare results across different platforms and algorithms.

Opportunity: Problem decomposition (PD) techniques are useful for addressing the problems that quantum computers encounter when computing large chemical systems. Specifically, we believe the use of PD techniques to decompose a target molecular system into smaller subsystems requiring fewer computational resources and implementing these approaches within quantum-computing-based frameworks will be an important direction towards achieving quantum advantages as we progress through and out of the NISQ-era. Some examples of PD techniques worth noting are the fragment molecular-orbital (FMO) method[1], the divide-and-conquer (DC) technique[2], and the density matrix embedding theory (DMET)[3].

Assessment: The use of PD techniques and divide-and-conquer techniques could address many of the listed problems above. It needs to be an integrated process that implements various techniques to fragment calculations at every level of abstraction and calculations. Here are step-by-step ideas on how they might be implemented:

1. **Chemical systems:** It is important to pick the chemical systems amendable to the application of fragmentation methods. These systems are, for example, organic molecules with multiple functional groups.
2. **Fragmentation methods:** The choice of a fragmentation method is important. Here are a few examples of such methods: the elongation method[4], the fragment molecular-orbital (FMO) method[1], the divide-and-conquer (DC) technique[2], and the density matrix embedding theory (DMET)[3].
3. **Active site decomposition methods:** In some cases, it is possible to further decompose calculations by using active space decomposition methods (ASD)[5], which use the density matrix renormalization group algorithms. The fragment wave functions are described by complete or restricted active-space wave functions. ASD alleviates the large computational cost of active-space methods by tailoring the wave function ansatz. Using ASD, as an example, a dimer wave function is compactly expressed as a linear combination of tensor products of monomer wave functions. Another more recently developed method is the localized active space self-consistent field (LAS-SCF) method[6] developed for efficient calculation of strongly-correlated systems for which we have already suggested a quantum-computing-based implementation[7].
4. **Circuit cutting:** The generated quantum circuits are likely still to be too large to compute on quantum circuits. The potential solution is to use the quantum circuit approach[8] to split a quantum circuit into sub-circuits.
5. **Alternative approaches:** If the above steps still do not achieve the goal then the alternative approaches should be considered. The fermionic Hamiltonians can be reformulated as spin Hamiltonians. Another potential technique is to use embedded approaches.

Among issues to be addressed include: how fragments are subsequently coupled to achieve accurate long-range interactions and choosing from and comparing different fragmentation schemes. Fragmentation methods can also be computationally expensive, particularly when many fragments are used. Overall, we believe fragmentation methods are a useful tool for studying large chemical systems, but their limitations should be considered carefully.

Timeliness or maturity: The steps in the above section need to be further developed and integrated into one larger pipeline. This will allow significant applications of quantum-computing to complex chemical systems within the next few years as quantum computing devices continue to improve.

Distributed quantum algorithms, such as our proposed quantum PD techniques, are a natural choice for use on forthcoming networked quantum computers. Making effective use of limited and potentially expensive connections between small quantum computers will be key to getting the most out of the emerging technology.

References: [1] K. Kitaura et al. *Chemical Physics Letters* 313.3-4, 1999. [2] M. Kobayashi and H. Nakai. *Linear-scaling techniques in computational chemistry and physics: methods and applications*, 2011. [3] N. C. Rubin. *arXiv preprint arXiv:1610.06910*, 2016. [4] A. Imamura et al. *The Journal of chemical physics* 95.7, 1991. [5] S. M. Parker and T. Shiozaki. *The Journal of chemical physics* 141.21, 2014. [6] M. R. Hermes and L. Gagliardi. *Journal of chemical theory and computation* 15.2, 2019. [7] M. Otten et al. *Journal of Chemical Theory and Computation* 18.12, 2022. [8] T. Ayral et al. *SN Computer Science* 2.3, 2021.

Machine learning for scalable quantum control

PI: Yuri Alexeev, Principal Project Specialist, Argonne National Laboratory, yuri@anl.gov
Co-PI: A. Baris Ozguler, Research Associate
Co-PI: Fred Chong, Professor, University of Chicago

Call: Position Paper for Workshop on Quantum Computing and Networking
Topic: Algorithms and Compilation

Challenge: Despite significant progress in quantum software development, the primary obstacle to achieving practical, functional quantum machines in the NISQ-era and beyond is the impact of noise and errors in quantum hardware. The main reason is various forms of electromagnetic noise, which cause decoherence that reduces coherent lifetimes and the fidelity of quantum logic operations when quantum devices are manipulated by faulty classical hardware. This limitation critically constrains the scope of useful computations that can be performed on quantum hardware, as measured by parameters such as circuit depth or quantum volume.

Improving the resilience of quantum hardware against noise and errors is crucial for pursuing commercially viable applications. One solution to this problem involves implementing low-level error suppression strategies inspired by the field of quantum control. This approach draws from classical control engineering, which is often used to stabilize unstable hardware but requires modification of fundamental concepts to successfully apply to the quantum domain.[1]

Another critical requirement for quantum processors to achieve practical utility is scaling up hardware for quantum information processing. A key challenge is the sustained precision control of large quantum systems. Many current quantum computer architectures consist of a layered control stack between algorithms and physical implementation. Scaling up requires optimization of pulses and communication between the layers. Efficient, accurate, and sometimes large-scale classical simulations of quantum devices provide a necessary glue between the layers, allowing for the practical generation of quantum optimal control pulses. An especially time-urgent challenge facing the development of quantum hardware is the need to reliably predict and control the pulses interacting with the device while maximizing computing power. This mission requires innovation in the development of improved data-driven and model-based approaches that maximize performance in existing experiments with an impact on optimizing operational scenarios. The development of a benchmark suite of data sets and corresponding machine learning (ML) models could lead to more rigorous algorithmic development and experimental studies on quantum simulators with direct interaction with hardware. The datasets and models would cover representative data sizes, diverse data models, data types, and formats. Such synthetic or experimental data may be provided by Q-NEXT, SQMS cavity-qubit device, microelectronics ASIC AI control team at Fermilab.

Despite the promising character of quantum devices, new classical simulation challenges arise. The general problem is that these systems are notoriously difficult to model and simulate using conventional techniques because the computational burden grows exponentially with the system size. Conventional algorithms for quantum control are computationally demanding, and thus limited to small-scale simulations. To find scalable solutions, simulators exploiting advanced HPC systems have been under active development. For example, there are quantum simulators developed at Argonne that scale on hundreds of supercomputer GPU nodes[2–5]. They can utilize high-speed interconnects, parallel file systems, nodes with specialized compute capabilities, several GPUs, and many cores per node to move data and direct computations. Detailed evaluations, tests, and integrations of these simulators into quantum systems are highly needed while both theory and experiment are evolving into a more mature state.

Opportunity: The following challenges and opportunities need to be addressed to advance the field of quantum computing:

1. Improve the performance of actual quantum hardware by providing optimized control strategies. This involves efficient characterization of error sources, identifying and utilizing system

controllability, and generating instructions for real hardware to mitigate the impact of noise and imperfections at the device level.

2. Provide increased functionality from limited quantum computational resources, as measured by parameters such as qubits, gates, and compute runtime. And make it accessible to users with varying levels of experience and expertise in quantum computing hardware or quantum control.
3. Ensure continued access to complex and rapidly advancing technology, as well as provide cutting-edge computational resources for tasks that require intensive numerical calculations via a modern cloud-compute architecture and supercomputers. This involves facilitating access to numerical techniques that benefit from or rely on specialized computational hardware such as GPUs.
4. Establish cross-compatibility with existing workflows, programming languages, QC architectures, and access methods. These tools should be seamlessly integrated into conventional programming workflows using Python, linking them to research code, cloud-based quantum computers, supercomputers, and custom QC hardware.

Assessment: To achieve the above-listed opportunities, hardware-software AI co-design approach is required to design robust gates that are resistant to cross-chip deviations. Autonomous learning prevents expensive new fabrication process especially when the physics model of the device is not clear. Modern reinforcement learning techniques can extract the device information using AI agents with high sample efficiency and with a linear scaling of the circuit depth[6]. AI agent interacts directly with a real quantum system and learns system response to the applied actions. It is desirable to apply control techniques in quantum-computational environments with real-time decision-making capability, necessitating purpose-built control hardware that can function in low-temperature environments. By approximating existing quantum control algorithms using ML, one can achieve comparable functionality using significantly less computation[7]. The resulting model could then be implemented as a hardware accelerator (ASIC/FPGA) at reasonable costs. Deep learning capability helps optimize device performance and capability in real-time. The scale of the data and model size will need to be significantly reduced to enable evaluations on the existing quantum circuit simulators and custom hardware (ASIC/FPGA) for quantum computers. As a part of the project, we will explore how large datasets can be reduced using ML dimensionality reduction techniques (feature selection, autoencoder etc.) to fit data on hardware. While some success is demonstrated on image data, representation learning for time series is a largely unexplored subject, which is of critical relevance to quantum control data consisting of rapidly changing signals. Time series representation learning techniques such as Contrastive Predictive Coding[8] can be adapted to use quantum resources, using methods such as computing representation similarities by evaluating the inner products of quantum states. In this framework, encoders are trained such that samples with the same context are similar, and those with different contexts should have low similarity.

Timeliness or maturity: The proposed HPC-driven ML simulations will guide hardware teams develop scalable quantum computer architectures. Supercomputers have always been crucial to drive developments in physics[9]. We expect a mutual benefit. Moving beyond the traditional supercomputer architecture utilizing only classical processors, quantum-centric supercomputers utilizing both quantum and classical processors will transform the way we solve scientific problems.

References: [1] H. Ball et al. *Quantum Science and Technology* 6.4, 2021. [2] D. Lykov et al. *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE. 2022. [3] D. Lykov et al. *2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS)*. IEEE. 2021. [4] W. Berquist et al. *arXiv preprint arXiv:2210.15874*, 2022. [5] X.-C. Wu et al. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 2019. [6] V. Sivak et al. *Physical Review X* 12.1, 2022. [7] D. Xu et al. *2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)*. IEEE. 2022. [8] A. v. d. Oord et al. *arXiv preprint arXiv:1807.03748*, 2018. [9] C. W. B. Davoudi et al. *arXiv preprint arXiv:2204.03381*, 2022.

Preparing for early fault-tolerant quantum computers

Andrew D. Baczewski (adbacze@sandia.gov), Alina Kononov (akonono@sandia.gov),
Shivesh Pathak (sapatha@sandia.gov), and Mohan Sarovar (mnsarov@sandia.gov)
Quantum Algorithms and Applications Collaboratory, Sandia National Laboratories

Topic: codesign and integration

Challenge: Many of the operations that are "easy" on noisy intermediate-scale quantum (NISQ) computers are likely to be "hard" on fault-tolerant quantum computers (FTQCs), and vice versa. For example, arbitrary single-qubit rotations are easy to implement on NISQ hardware but require relatively expensive rotation synthesis [6] and magic state distillation [2] on FTQCs. CNOTs are a limiting factor in the NISQ era but relatively straightforward to implement fault-tolerantly on an FTQC [3, 4]. As a consequence, codesign strategies for efficient NISQ algorithms will lead them to be especially poorly suited to early FTQCs, in which access to high fidelity resource states will be a limiting factor. Short circuits with many concurrent single-qubit rotations - ideal circuits in the NISQ era - will become particularly resource intensive due to the need to dedicate a sizable fraction of a machine to preparing and injecting resource states. For example, a simple "2-qubit" problem like solving the ground state of molecular hydrogen using the variational quantum eigensolver (VQE) would require 15 logical qubits in a surface code architecture using lattice surgery and a single T factory in a somewhat naive implementation [4]. This overhead is compounded as more arbitrary rotations, and thus either more factories or more idling, are needed. Simply put, many of the NISQ-era lessons that we've learned about algorithm design, compilation, and overall performance optimization, are likely to lead us astray as early FTQCs become available.

Opportunity: Practical algorithm development will need to adapt to both new architectural constraints and architectural constraints that are in direct opposition to the ones that we've been optimizing for in the NISQ era. NISQ-era benchmarks will also need to be adapted or replaced, as simple measures of computational volume will fail to reflect the real quality and capability of an early FTQC. Programming models and compilers will also need to become increasingly conscientious of the new set of trade-offs that these machines will bring with them.

Assessment: Success could involve identifying common codesign strategies that allow us to seamlessly transition from the NISQ era to the early FTQC era and beyond. For example, we should be prepared to answer questions about what we would do with machines with hundreds or thousands of qubits with error rates steadily working below thresholds for quantum error correction (QEC). What are the likely limits of error mitigation and viable circuit volumes without resorting to QEC, and what will we be able to do on FTQCs that are just on the other side of this boundary? The successful design of credible applications for early FTQCs - even simple "Hello world" demonstrations like logical VQE - will explicitly articulate a path forward for whenever we reach the end of the NISQ era. It is also possible that the transition isn't seamless, in which case clearly articulating the challenges to come will be valuable in and of itself.

Timeliness or maturity: QEC research is outside of ASCR's mandate, but it is timely to support research that is at least aware of its consequences for all of the topics that are within ASCR's scope. Recent demonstrations [1, 5] indicate that it is likely that QEC and increasingly sophisticated demonstrations of FTQC will become viable on the testbeds that ASCR can access. If we don't anticipate this literal phase transition, then many of the tools and techniques that we have invested in will be rapidly deprecated. To further emphasize timeliness, consider the implementation of the “2-qubit” logical VQE problem mentioned above. A naive minimal implementation could be implemented on 255 sufficiently performant physical qubits (i.e., distance-3 surface code). This qubit count is conspicuously smaller than IBM’s largest machine, but the quality of those qubits is almost certainly insufficient to implement it. Further, the first fully logical VQE demonstration is likely to be more bespoke than this naive picture, but it will likely be on a machine with hundreds to thousands of physical qubits. It does not seem unreasonable to imagine machines capable enough to implement such an exemplar becoming available in the next 5 years and ASCR should be prepared to adapt.

References

[1] Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949):676–681, 2023.

[2] S. Bravyi and A. Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Physical Review A*, 71(2):022316, 2005.

[3] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter. Surface code quantum computing by lattice surgery. *New Journal of Physics*, 14(12):123011, 2012.

[4] D. Litinski. A game of surface codes: Large-scale quantum computing with lattice surgery. *Quantum*, 3:128, 2019.

[5] C. Ryan-Anderson, N. Brown, M. Allman, B. Arkin, G. Asa-Attuah, C. Baldwin, J. Berg, J. Bohnet, S. Braxton, N. Burdick, et al. Implementing fault-tolerant entangling gates on the five-qubit code and the color code. *arXiv preprint arXiv:2208.01863*, 2022.

[6] P. Selinger. Efficient clifford+ t approximation of single-qubit operators. *arXiv preprint arXiv:1212.6253*, 2012.

Machine Learning Augmented Quantum Computing

Avah Banerjee
Missouri University of Science and Technology
Email: banerjeeav@mst.edu

Introduction. In recent years, there has been a rapid growth in interest in using variational quantum algorithm (VQA) [1] models to augment classical machine learning. This surge in research has become one of the main drivers in the field of quantum machine learning (QML), particularly focused on leveraging near-term quantum hardware. Despite this, there have been limited theoretical insights into the power and limitations of (variational) QML compared to its classical counterparts, and experimental evidence is constrained by the inability to achieve scalability with limited-qubit pre-fault tolerant (PFT)[1] hardware. In this context, we explore the opposite direction – how to accelerate quantum computing by leveraging the rapidly growing capabilities of classical machine learning tools. Our central hypothesis is that there may be ways to design quantum programs using ML techniques in an efficient manner, even if the quantum processes realized by these programs are fundamentally hard to simulate using purely classical means.

Challenge. We broadly examine two distinct domains, and for each, we identify several high-level challenges. **(C1) Quantum Compilation**: Quantum compilation is the high-level process that begins with a quantum program (written using a quantum software toolkit) and converts it into appropriate hardware instructions for a given quantum hardware. This process consists of an extensive and interdependent pipeline with numerous components, including circuit synthesis and optimization, transpilation, hardware-specific qubit mapping and routing, gate scheduling, and optimization of low-level laser pulse controls, scheduling of of trapped-ion transport etc. Building a robust quantum program-to-circuit execution pipeline requires careful optimization of each stage in the pipeline. For example, the circuit generated after adding in the necessary gates to satisfy connectivity constraints of the quantum processor will affect how we schedule the gate operations. The latter will affect the noise due to cross-talk between adjacent two-qubit gates. Despite these interdependencies, it is computationally infeasible to perform a global optimization of all stages simultaneously; a modular approach is usually prescribed. These approaches vary between runtime systems and specific hardware architectures. In summary, the quantum compilation problem, as it currently stands, is optimized in an ad-hoc manner, and the challenge lies in moving beyond this towards a system-independent universal framework.

(C2) Robust Quantum Simulation: Efficient Simulation of quantum system on PFT quantum computers is a challenging task. We highlight two important problems in this area. *A. Generation of wavefunction ansatz for VQAs.* In the VQA framework, a circuit is generated from some wavefunction ansatz for the target wavefunction (e.g., for the ground state energy estimation problem). Despite recent breakthroughs (e.g., ADAPT-VQE), constructing parsimonious ansatz circuits without increasing number of measurements remains a challenge. *B. Simulating a given Hamiltonian.* Although several asymptotically optimal algorithms have been developed in recent years (quantum signal processing, qubitization etc.), their additional resource requirements, such as ancilla qubits, make them impractical for PFT hardware and product formula-based approaches (particularly those leveraging randomization - e.g. QDRIFT or its recent higher order generalizations) still achieve better empirical performance. Despite this, the circuits needed to simulate all but the most simplest of electronic hamiltonians exceeds the quantum volume of PFT hardware by several orders of magnitude [2].

[1]This term is employed in place of NISQ, as it encompasses a wider range of technologies, including both near-term and future developments, that lack scalable error correction.

Opportunity. **(C1)** Using machine learning techniques to produce robust quantum circuits is an underexplored area of research. On the other hand ML based code generation tools for classical computers have become increasingly widespread in the past few years. Even though it is possible to use transformer-type NNs to produce simple high level quantum programs using standard programming languages, it is a much harder task implement produce non-trivial circuits optimized for a target hardware. However, graphical structure of quantum circuits makes them suitable objects for Graph Neural Network based ML architectures which has already seen usage in conventional circuit design. It is also possible given a specific quantum hardware to collect enough data to devise hardware specific compilation strategies automatically using ML tools.

(C2) In classical physics, specialized neural network models, such as Hamiltonian neural networks and Lagrangian neural networks, have been shown to approximate the underlying physical laws of simple systems effectively. Neural network models have also been developed to approximate quantum states (e.g., neural network quantum states) and recognize certain quantum contextuality (e.g., CHSH-type). It is crucial to investigate the extent to which we can model quantum system Hamiltonians as well. Neural networks might be employed to construct Hamiltonians with fewer terms than traditional transformation methods. Furthermore, it could even be feasible to directly convert a fermionic Hamiltonian into a quantum circuit over a given gate set, bypassing several steps in the Hamiltonian simulation pipeline. Finally, in variational quantum algorithms (VQAs), the primary challenge lies in optimizing the ansatz parameters [2]. However, if the ansatz structure is not trainable for a specific problem, it could result in suboptimal outcomes. It may be possible to guide the ansatz construction process using machine learning (ML) – either through automatic generation (e.g., diffusion models on graphs) or sequential augmentation (e.g., using reinforcement learning).

Assessment. **(C1)** A wide range of publicly available quantum software development kits and packages exists today. The majority of these offer program-to-circuit pipelines, complete with hardware-specific optimizations. Notably, some of these systems enable users to substitute existing pipeline blocks with custom methods. Automated pipelines developed using ML-based approaches can be benchmarked against these existing toolkits, for either partial or complete stages of the compilation process. Robustness to hardware noise can be evaluated through simulation or by running on actual hardware. Metrics such as gate count, circuit depth, and output state fidelity can be utilized to compare various schemes. **(C2)** Automatically constructed VQA ansatzes can be assessed based on their trainability, which includes the ability to avoid barren plateaus and maintain robustness against hardware noise, among other factors. Similarly, the accuracy of automatically compiled circuits for a given Hamiltonian can be evaluated by examining the precision of the estimated parameters (e.g., ground state energy).

Timeliness. (1) Recent advancements in machine learning tools, particularly generative models, have demonstrated great potential in generating and simulating biological, chemical, and physical structures and processes. We believe that these models can be extended to the quantum setting, potentially enhancing the capabilities of PFT quantum computers. (2) With the advent of cloud-based quantum computers, the adoption of quantum computing has significantly accelerated. Consequently, generating training data for the proposed tasks may prove to be easier than what was possible just a few years ago, supporting further progress in this domain.

References

[1] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021.

[2] Andrew M Childs, Dmitri Maslov, Yunseong Nam, Neil J Ross, and Yuan Su. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38):9456–9461, 2018.

[2] There have already been studies on using NN's to optimize VQA parameters.

Integrating Quantum Co Processors with High Performance Computers

Kristin Beck, Yaniv Rosen, Jonathan DuBois
Lawrence Livermore National Laboratory, Livermore, CA 94550
{beck37, rosen10, dubois9}@llnl.gov

Topic: Integration

Quantum computers are anticipated to vastly improve computational efficiency for certain problems. We envision a future where quantum processing units are part of the ecosystem of a traditional high performance computing (HPC) environment and accessed as a resource during computation. HPC environments feature parallel computing with computing resources linked via high-speed, low-latency connections [1]. The integration models could include quantum subroutines such as order finding [2] and iterations of variational algorithms [3] or queries of a pre-prepared quantum state that simulates a subsystem that may be difficult to calculate. DOE's ASCR should anticipate the integration of quantum coprocessors with HPC and develop HPC integration testbeds or add HPC integration into existing testbed. These testbeds should include quantum hardware, classical coprocessors, and small HPC environments such as few-node clusters to build tools that enable practically-motivated exploration of the new tools needed for scheduling, compilation and execution [4]. They will illuminate where current models for integration fail and generate refined ideas for system topology that anticipate and accelerate the integration of future quantum hardware, including long coherence time hardware with partial measurements where HPC results may be integrated into ongoing quantum calculations.

While DOE has long been interested in how to work efficiently in a heterogeneous computing environment that incorporates novel computing technologies [5], the reduction to practice with quantum computers so far has been primarily in the private sector or with typically high-latency cloud-integrated technologies. Working with physical hardware also enables full transparency into the characterization and control problem that plagues quantum devices, which often require fine tuning for every qubit and pair of qubits in parameters such as frequencies, coupling strengths, or bias voltages to get decent performance of the quantum logical operations.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-CONF-848505

References

1. “What is HPC? Introduction to High Performance Computing” ibm.com/topics/hpc. Accessed 25 April 2023.
2. P. Shor “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer.” SIAM Review 41 (2) 1999
   Nielsen, Michael A. & Chuang, Isaac L. “Quantum Computation and Quantum Information, 10th Anniversary Edition” Cambridge University Press, ISBN 9781107002173. 2010
3. Peruzzo, A. *et al*. “A variational eigenvalue solver on a photonic quantum processor.” Nat. Commun. 5 (4213) 2014
   Y. Nam *et al.* “Ground-state energy estimation of the water molecule on a trapped ion quantum computer.” npj Quantum Information 6 (33) 2020
4. T.S. Humble *et al.* “Quantum Computers for High Performance Computing.” IEEE Micro 41 (5) 2021.
5. A. Aspuru-Guzik *et al.*, "ASCR Workshop on Quantum Computing for Science", 2015, https://www.osti.gov/biblio/1194404.
   J. S. Vetter *et al.*, "Productive Computational Science in the Era of Extreme Heterogeneity: Report for DOE ASCR Workshop on Extreme Heterogeneity", 2018, https://www.osti.gov/biblio/1473756.

Variational Quantum Computing in the NISQ Era and Beyond

Ryan Bennink[1], Sarah Chehade, Pras Date, Kathleen Hamilton, Phil Lotshaw, Titus Morris, Francisco Rios, Joseph Wang, Yan Wang

Oak Ridge National Laboratory

Topic: Quantum algorithms

Challenge: Fault-tolerant application-scale quantum computing appears to be many years away. In the meantime, variational quantum computing (VQC)—in which the parameters of a tunable quantum circuit are optimized at run time to achieve the desired result—has emerged as the most practical approach to quantum computing in the current and foreseeable era of noisy, intermediate-scale quantum (NISQ) devices. While VQC algorithms have been demonstrated to solve very small problems in a wide variety of domains including chemistry, optimization, and machine learning, *the prospects for VQC to provide utility at scale are uncertain*. Recent theoretical works indicate that the circuit ansatze employed in many existing VQC algorithms will become untrainable [1,2] and produce inferior solutions [3–5] at application scale. It is possible that carefully designed problem-specific VQC methods may overcome such limitations and provide computational utility at scale [6–8], but there is currently insufficient understanding of the interplay between problem type, ansatz structure, objective function, and optimization strategy to make such algorithmic advances.

Opportunity: The quantum computing community has an opportunity to develop new variational quantum algorithms that provide utility deep into the NISQ era and beyond—*or* to prove conclusively that such a dream can never be realized. To obtain either outcome will require continued fundamental research in the theory of parametric quantum circuits. Such research will either inform new approaches to the design of NISQ-friendly algorithms that are *by construction* trainable, noise-robust, and effective at scale, or show why such algorithms cannot exist.

Quantum Circuit Theory Research. Recent works have begun to develop the theory of parametric quantum circuits in the context of VQC, introducing mathematical tools and perspectives to investigate issues such as barren plateaus [1,2], spurious local optima [9], expressiveness [10], optimization complexity [11], overparameterization [12], and symmetries [13]. Efforts of this kind must continue and expand in order to yield rigorous results and new insights into the design and capabilities of parametric quantum circuits. Some key questions to be addressed are the following: *(i)* What is the relationship between circuit structure and the states or unitaries that can be implemented? *(ii)* How does the choice of loss function impact the loss landscape, and hence trainability, of the circuit? *(iii)* Are there practically relevant quantum classes (perhaps encoding classical models) that are difficult to learn or compute using classical computers, that can be efficiently learned or computed using quantum circuits? *(iv)* In kernel methods [14], what is relationship between the data-to-circuit mapping, observable, and power of the kernel function? Are there powerful kernel functions that can be computed more efficiently using quantum circuits? *(v)* Which phenomena and techniques from classical variational methods carry over to quantum variational methods? *(vi)* How can unitarity and entanglement, which are not present in classical variational models, be leveraged for advantage?

Algorithms Research. The first generation of variational quantum algorithms—including the variational quantum eigensolver (VQE) and quantum approximate optimization algorithm (QAOA)—were developed without the benefit of any well-developed theory of parametric quantum circuits. It is now

[1]Email: benninkrs@ornl.gov

known that these algorithms face strong barriers to scalability. However, there are indications that new approaches informed by better theoretical understanding of parametric circuits may not face such barriers. Promising directions include adaptively grown [8] or feedback-driven ansatze [15], the use of nonlinear cost functions [6], incorporation of domain symmetries [13,16,17], and smart parameter initialization [18]. Other possibilities for developments include incorporation of non-unitary operations to reduce the impact of noise, methods to discover symmetries in data sets, and the use of novel scalable and noise-resilient classical optimization methods.

Assessment: Research activities along the lines described above could have several kinds of successes. The first kind of success would be a set of rigorous analytical bounds on the performance of representative VQC methods as functions of problem scale and implementation error. Another kind of success would be several new ansatz-based quantum algorithms with known performance and cost scaling for well-defined tasks in several different domains.

Timeliness: Thanks to the recent advent of publicly accessible quantum processors with dozens of qubits, the ability of quantum computers to solve a wide variety of problems at trivially small scales has by now been well-established. A key question to be answered in the next 3-5 years is whether any utility can be found in quantum computing prior to the hoped-for advent of huge fault-tolerant quantum processors many years from now. A deeper theory of parametric quantum circuits that informs the development of a new generation of scalable, noise-resilient variational quantum algorithms will go a long way toward answering this key question and will shape quantum computing research for decades.

References

[1] J. R. McClean *et al.*, Nat. Commun. **9**, 4812 (2018).
[2] Z. Holmes *et al.*, PRX Quantum **3**, 010313 (2022).
[3] D. Stilck França and R. García-Patrón, Nat. Phys. **17**, 1221 (2021).
[4] G. González-García, R. Trivedi, and J. I. Cirac, PRX Quantum **3**, 040326 (2022).
[5] G. De Palma *et al.*, PRX Quantum **4**, 010309 (2023).
[6] M. Kieferova, O. M. Carlos, and N. Wiebe, (2021).
[7] A. Pesah *et al.*, Phys. Rev. X **11**, 041011 (2021).
[8] H. R. Grimsley *et al.*, Npj Quantum Inf. **9**, 1 (2023).
[9] E. R. Anschuetz and B. T. Kiani, Nat. Commun. **13**, 7760 (2022).
[10] Y. Du *et al.*, Phys. Rev. Lett. **128**, 080506 (2022).
[11] L. Bittel and M. Kliesch, Phys. Rev. Lett. **127**, 120502 (2021).
[12] M. Larocca *et al.*, ArXiv210911676 Quant-Ph Stat (2021).
[13] J. J. Meyer *et al.*, PRX Quantum **4**, 010328 (2023).
[14] M. Schuld, ArXiv210111020 Quant-Ph Stat (2021).
[15] A. B. Magann *et al.*, Phys. Rev. Lett. **129**, 250502 (2022).
[16] B. T. Gard *et al.*, Npj Quantum Inf. **6**, 10 (2020).
[17] F. Sauvage *et al.*, arXiv:2207.14413 (2022).
[18] K. Mitarai *et al.*, Phys. Rev. Res. **4**, 033012 (2022).

Accurate hierarchical models of quantum computers are critical for assessing performance

Robin Blume-Kohout (*rjblume@sandia.gov*), Kevin Young (*kyoung@sandia.gov*), Megan Dahlhauser

Quantum Performance Laboratory, Sandia National Laboratories

Topic: models, error correction and mitigation, and codesign and integration

Challenge: The models and modeling of integrated quantum computing processors that are available today are catastrophically inaccurate [1], and will fail to support and enable useful quantum computing. Predictive modeling plays several essential roles in the ongoing quest for useful quantum computing, which is expected to require (at a minimum) 100,000 – 1,000,000 physical qubits with gate error rates of 10^{-3} – 10^{-4} [2,3]. Many engineering advances will be required, since the current state-of-the-art is ~100 qubits with 1% error per 2-qubit gate. Predictive modeling of device behavior and errors is necessary to:

- Evaluate, choose, and refine architectures that support algorithms and error correction,
- Design good error correcting codes and decoders,
- Identify design bottlenecks and opportunities, to allocate R&D resources wisely,
- Evaluate, verify, and validate prototype devices using characterization/benchmarking data.

These tasks are essential to rapid and efficient R&D progress toward useful hardware. They require *predictive* models of qubits, gates, and integrated quantum devices that can accurately predict how observed fault mechanisms or subtle changes to quantum components will impact scaled up next-generation processors. These models must predict the outcomes of quantum circuits, and they must do so at the same accuracy that is required for practical fault tolerant quantum error correction (FTQEC) [4], which is expected to be about 10^{-4} per gate [2,3].

But current state of the art models do not predict with this accuracy – they are off by 50x or more! Our group evaluated 12 leading testbeds in 2019 [1], and found that "standard error metrics are poor predictors of whether a program will run successfully" (see Figure 1). We found large discrepancies between experimental results of running quantum circuits, and the predictions of the manufacturers' published error rates. Very recently (Figure 2), our group extended this study quantitatively by focusing on just 2 qubits of an IBM Q device, and evaluating the predictions of both an IBM-provided emulator *and* our own cutting-edge *gate set tomography* [5] model. Both models made predictions that deviated by at least 0.5% / gate for some circuits – 50x higher than the accuracy expected to be necessary for enabling FTQEC. **We conclude that because currently-available models do not capture real-world behavior of quantum computers, breakthroughs in faithfulness and scalability are badly needed to enable useful quantum computing.** We suggest that DOE/ASCR can help meet this urgent need by encouraging development of cross-disciplinary hierarchical models and modeling techniques for quantum computing *components* (qubits, gates) and integrated quantum computing *testbeds*.

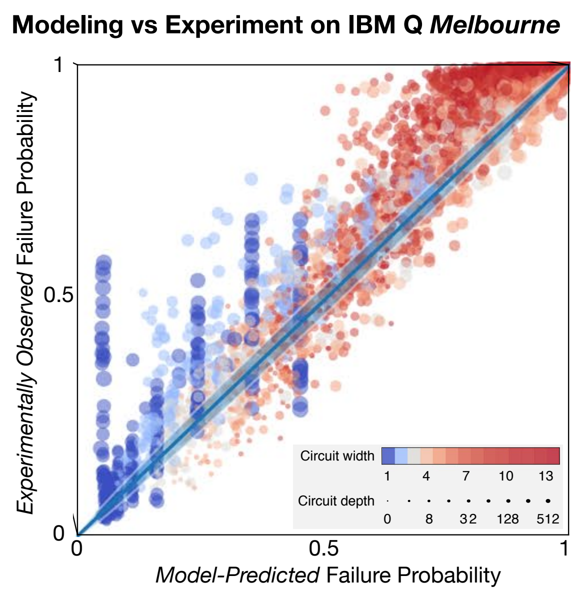

Figure 1: We tested the predictive accuracy of manufacturer-provided error rates by running a wide variety of quantum circuits on IBM Q *Melbourne* in 2019. Most results (including both wide and deep circuits) deviated by >2σ (indicated by blue band along diagonal) from model predictions. (Reproduced from Ref. [1].)

Opportunity: Accurate, scalable models of errors in quantum computers is an outstanding scientific challenge. We identify three *specific* challenges, which need to be addressed in parallel and (eventually) brought together to solve the problem. **First**, the detailed device physics of common qubits (transmons [6], electrons in semiconductors [7], trapped atoms or ions [8]) leads to complex non-Markovian [9,10] behavior that *must* be captured accurately before it can be propagated up the stack. **Second**, scalable modeling of *N* qubits demand more efficient representations than the current gold standard of process matrices, which grow exponentially with *N* [11]. **Third**, integrated quantum processors (including not just qubits but also control systems) display complex emergent behaviors that must be captured in order to predict qubit behavior accurately. Each of these challenges requires focused multidisciplinary teaming and careful quantitative assessment (e.g. by constantly monitoring and improving predictive accuracy). The unique opportunity for DOE/ASCR is to encourage and support the creation of such teams, *beyond* the single-investigator level, providing the long-term vision for accurate predictive modeling that will enable viable quantum computing.

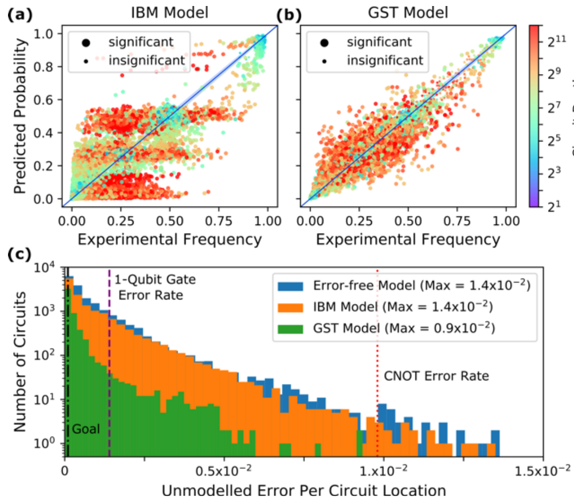

Figure 2: We ran 25,000 *gate set tomography* circuits on 2 qubits of an IBM Q device in 2023, and compared data to **(a)** predictions by IBM's own emulator and **(b)** a gate set model fit to this specific data. Neither model captures reality, and examination of both models' per-gate prediction error **(c)** shows 0.5%-1%/gate inaccuracy.

Assessment: What is needed here – and therefore the criterion for success – is simple: modeling is sufficiently accurate and scalable *if* it enables designers of hardware, architecture, and algorithms to confidently answer "*What if* we did X?" by simulation instead of experiment. Experiment will always be necessary to evaluate transformative unmodelable ideas, but right now even quotidian hypotheses ("What if errors were a little different?") require experiment [12]. Simulation must replace experiment for these questions, and this defines what we need from models. There is a simple proxy for this: models of *existing* hardware must be able to predict circuit outcomes to within 10^{-4} total variation per gate.

Timeliness or maturity: The urgency of this challenge is obvious: qubit hardware outpaced modeling almost 10 years ago [13]. Quantum computing has already transitioned to an engineering-driven field, and the engineering is critically handicapped by the total absence of accurate modeling. The cutting edge of qubit performance must transition from 0.5-1% gate error rates to 10^{-4}-10^{-3} gate error rates. In this transition, accurate and scalable modeling will become absolutely essential. If we cannot *predict* integrated quantum devices, we do not *understand* them – and without deep, scalable understanding we cannot design and engineer next-generation and future quantum computing devices.

References

[1] T. Proctor, K. Rudinger, K. Young, E. Nielsen and R. Blume-Kohout, "Measuring the Capabilities of Quantum Computers". *Nature Physics* **18**, 75-79 (2022).
[2] D. Wecker *et al*. "Gate-count estimates for performing quantum chemistry on small quantum computers". *Physical Review A* **90**, 022305 (2014).
[3] C. Gidney, and M. Ekerå, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits." *Quantum* **5**, 433 (2021)
[4] A. Calderbank and P. Shor, "Good quantum error-correcting codes exist." *Physical Review A* **54**(2), 1098 (1996)
[5] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young and R. Blume-Kohout, "Gate set tomography". *Quantum* **5**, 557 (2021).
[6] M. Kjaergaard *et al.*, "Superconducting Qubits: Current State of Play," *Annu. Rev. Condens. Matter Phys.* **11**(1), 369–395 (2020).
[7] M. Mądzik *et al*. "Precision tomography of a three-qubit donor quantum processor in silicon." *Nature* **601**, 348-353 (2022)
[8] L. Postler *et al*, "Demonstration of fault-tolerant universal quantum gate operations". *Nature* **605**, 675-680 (2022).
[9] F. A. Pollock et al. "Non-Markovian quantum processes: Complete framework and efficient characterization." *Physical Review A* **97**(1), 012127 (2018).
[10] B. M. Terhal, Barbara and G. Burkard. "Fault-tolerant quantum computation for local non-Markovian noise." *Physical Review A* **71**(1), 012336 (2005).
[11] R. Blume-Kohout *et al*, "A taxonomy of small Markovian errors". *PRX Quantum* **3**(2), 020335 (2022)
[12] P. Iyer and D. Poulin "A small quantum computer is needed to optimize fault-tolerant protocols." *Quantum Science and Technology* **3**(3), 030504 (2018).
[13] R. Barends *et al*. "Superconducting quantum circuits at the surface code threshold for fault tolerance." *Nature* **508**(7497), 500-503 (2014).

Computational fluid dynamics on quantum platforms: Current and future directions

Spencer H. Bryngelson,[1] Bryan Gard,[2] and Alexander Alexeev[3]
[1] School of Computational Science & Engineering,
Georgia Institute of Technology, Atlanta, GA 30332, USA*
[2] CIPHER Laboratory, Georgia Tech Research Institute, Atlanta, GA 30332, USA
[3] George W. Woodruff School of Mechanical Engineering,
Georgia Institute of Technology, Atlanta, GA 30332, USA

Topic: Quantum algorithms for fluid dynamics problems

Computational fluid dynamics (CFD) has driven the development of numerical methods and algorithms for classical hardware for the last 70 years, starting with computations on the ENIAC machine orchestrated by Jon von Neumann [1]. Problems in fluid dynamics are characterized by the solution to the Navier–Stokes equations, which are a mixed system of nonlinear partial differential equations (PDEs) with hyperbolic terms associated with fluid inertia and parabolic ones associated with viscosity and mass conservation. Work on solving these problems using numerical methods is ongoing for at least the past 100 years [2] For example, the finite difference, volume, and element techniques have roots in solving Navier–Stokes-like equations. CFD computations make up a significant fraction of supercomputer use in the United States, particularly on leadership-class systems maintained by the US Department of Energy, with application to weather and climate forecasting, air, land, and sea vehicle design, air and hydraulic energy harvesting, and more. These computations are expensive, though. For example, the cost of performing a scale-resolving simulation of turbulence increases with Reynolds number (ratio of inertial to viscous effects) as $\mathcal{O}(\mathrm{Re}^3)$.

Challenge: Continuum Mechanics and PDEs

Quantum computers make exponential speedups possible in specific cases of carefully crafted algorithms (viz. Grover and Shor). In the natural sciences, quantum computers are most closely associated with quantum-scale problems, like quantum chemistry [3]. This perspective is shared. For example, in Richard Feynman's 1981 lecture *Simulating Physics with Computers*, he stated, "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy." Quantum-scale CFD simulations are impractical for continuum-scale problems, partly due to algorithmic constraints like the CFL condition. Thus, our hand appears forced to address continuum-scale CFD problems at their own scale, solving the corresponding PDEs, the Navier–Stokes equations

$$\frac{\partial \boldsymbol{u}}{\partial t} + (\boldsymbol{u}\cdot\nabla_{\boldsymbol{x}})\boldsymbol{u} = -\nabla_{\boldsymbol{x}} p + \frac{1}{\mathrm{Re}}\nabla^2_{\boldsymbol{x}}\boldsymbol{u} \qquad \text{with} \qquad \frac{\partial \rho}{\partial t} + \nabla_{\boldsymbol{x}}\cdot(\rho\boldsymbol{u}) = 0, \tag{1}$$

where $\boldsymbol{u}(\boldsymbol{x},t)$ and $\rho(\boldsymbol{x},t)$ are velocity and density fields (dependent variables) and $\boldsymbol{x}$ and t are space-time coordinates.

There are many ways to solve (1) and its variants, each with its own relative merits. Numerical schemes range from microscopic methods based on particles, like dissipative or smoothed particle dynamics, to mesoscale methods that use kinetic closures like the lattice Boltzmann method, to macroscopic methods like finite volume, difference, and element schemes. However, analysis of each method's efficiency, complexity, and robustness is understood in terms of its implementation on classical hardware. If we hope to solve (1) efficiently on *quantum hardware*, then we must appreciate the trade-offs associated with quantum algorithm surrogates of established techniques. It is further unclear if quantum analogs of classical algorithms for nonlinear mixed PDEs are appropriate for quantum computers or if one should construct new methods based on established or supposed quantum advantage.

Opportunity: Hybrid Algorithms for the Navier–Stokes equations

The landmark work of Harrow *et al.* [4] established the first quantum algorithm, HHL, for solving linear systems of equations with exponential speedup. Such improvements can be marshaled to solve part of (1) via the method of lines. Andrew Childs and coworkers have also developed quantum finite difference and spectral methods for differential equations with provable complexity improvements [5]. However, these methods require deep many-gate circuits to

* shb@gatech.edu

achieve results comparable to what one can obtain on even a modern laptop. This proves prohibitive for current hardware but paints a bright future for larger, fault-tolerant quantum devices.

In the near term, variational quantum algorithms appear better suited for current NISQ (noisy intermediate-scale quantum) devices. For example, hybrid quantum–classical procedures can evaluate the solution quality via a cost function using a quantum computer and optimize variational parameters using a classical computer. Thus, variational methods enable quantum algorithms with relatively shallow gate depths and qubit counts. Bravo-Prieto *et al.* [6] introduced the variational quantum linear solver based on Hadamard tests, and Liu *et al.* [7] showed that a Poisson problem could be solved via the Quantum Alternating Operator Ansatz (QAOA) algorithm. However, work is still needed to clarify these techniques' improvements. For example, variation algorithms require an (at present) unclear ansatz and do not guarantee convergence or asymptotic speedups.

The continued tension between long- and short-term quantum algorithm threads has contributed to the welcomed improvement of each. However, for the foreseeable future, hybrid quantum–classical methods of some kind will be required to solve the large 3D problems that CFD entails.

Assessment: Speedups, Scaling, and Networking

The linchpin to solving CFD problems of practical interest is processing the large numbers of spatial grid points (or particles) and the commensurate large number of time steps for integration. For this, quantum computers have already established efficiency via near-term algorithms for linear problems, like the pressure Poisson equation associated with enforcing incompressibility constraints. What remains to be seen is the robustness and complexity of these methods for the required grid resolution of physically meaningful problems. Efforts in this area should focus on evaluating the ability of quantum hardware to realize accurate solutions to large Poisson-like equations and appropriate discretizations of them, including high-order accurate differencing and spectral schemes. To begin augmenting classical algorithms, quantum hardware should be able to solve such equations for at least millions of degrees of freedom. Problems of this size correspond to a reasonable number of qubits (less than 100), though the quantum processing time should be on the sub-second scale at near single precision accuracy (about 10^{-8}).

Quantum computers must network to classical ones for hybrid algorithms to be of practical utility. Considering the problem sizes involved, this means transferring state data at least at current data center-class bandwidth, around 100 GB/s. Quantum algorithms must process this data into qubit states at similar rates.

Timeliness: The Pressure of Moore's Law

As Moore's law struggles to keep pace, we are pushed into an ever less stable corner for the large-scale PDE problems associated with CFD. This trajectory has prompted efforts from NASA, like the CFD Vision 2030 report [8], to establish methods on "revolutionary platforms" to continue simulating and designing high-speed air and spacecraft. In 2014, at the time of the vision report, it was unclear what algorithms we could use or what platforms we should embrace. Today, we marshal knowledge of quantum algorithms that achieve *de facto* complexity improvements on even simplified CFD problems. Much work remains, but the groundwork has been set.

[1] J. G. Charney, R. Fjörtoft, and J. von Neumann, Numerical integration of the barotropic vorticity equation, Tellus **2**, 237 (1950).

[2] L. F. Richardson, *Weather Prediction by Numerical Process* (Cambridge University Press, 1922).

[3] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Quantum chemistry in the age of quantum computing, Chemical Rreviews **119**, 10856 (2019).

[4] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for solving linear systems of equations, Physical Review Letters **103**, 150502 (2009).

[5] A. M. Childs, J.-P. Liu, and A. Ostrander, High-precision quantum algorithms for partial differential equations, Quantum **5**, 574 (2021).

[6] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, Variational quantum linear solver, arXiv preprint arXiv:1909.05820 (2020).

[7] H.-L. Liu, Y.-S. Wu, L.-C. Wan, S.-J. Pan, S.-J. Qin, F. Gao, and Q.-Y. Wen, Variational quantum algorithm for the Poisson equation, Physical Review A **104**, 022418 (2021).

[8] J. P. Slotnick, A. Khodadoust, J. Alonso, D. Darmofal, W. Gropp, E. Lurie, and D. J. Mavriplis, *CFD vision 2030 study: A path to revolutionary computational aerosciences*, Tech. Rep. NF1676L-18332 (NASA, 2014).

Entanglement-based Networks: Towards Connection-oriented design?

Angela Sara Cacciapuoti, *Senior Member, IEEE*, Jessica Illiano, Marcello Caleffi, *Senior Member, IEEE*

I. Topic

Quantum network models and Architectures

II. Challenge

The Quantum Internet [1]–[5] has the potential of enabling applications with no counterpart in the classical Internet [6]. Indeed, the Quantum Internet, by generating and distributing entangled quantum states, comes with a whole new dazzling functionalities [5], [7], [8].

From a design point of view, entanglement enables a new and wider form of connectivity, with respect to classical networks [1]. Specifically, as long as an entangled state – say an EPR pair for the sake of simplicity – is shared between two nodes, a qubit can be "transmitted" regardless of the instantaneous conditions of the underlying physical quantum channel, via quantum teleportation. Remarkably, the qubit transmission is still possible even if the nodes are not anymore interconnected by a quantum link. In this sense, entanglement enables a new form of connectivity, referred to as entanglement-based connectivity. Furthermore, entanglement can be swapped and, hence, it is possible to dynamically – namely, at run time – change the identities of the entangled nodes.

As represented in Figure 1, when it comes to multipartite entanglement, the dynamic nature of the entanglement-based connectivity becomes even more evident. As instance, by distributing an n-qubit GHZ state among n network nodes, an EPR pair can be distributively extracted by any pair of nodes, with the identities of the entangled nodes chosen at run-time.

The entanglement-based connectivity has a profound effect on the entire Quantum Internet protocol stack, that should be grasped by an effective design. As instance, it redefines the same concept of neighborhood with no counterpart in the classical networks. Furthermore, it must be noted that entanglement constitutes an highly heterogeneous resource: there exist different classes of multipartite entangled states, which exhibit different properties and enable different network applications. And, such an heterogeneity has not been explored so far.

In this context, it is clear that an additional crucial challenge for quantum network design is constituted by the entanglement distribution process. Indeed, in order to exploit the entanglement-based connectivity the network nodes should be provided with the entanglement resources.

Besides, the strategy adopted for the distribution of the entanglement impacts on the connectivity. Specifically, there exist two different strategies for the entanglement distribution from a network engineering prospective: *proactive* or *reactive*. Proactive strategies aim at early distribution of entanglement resources – ideally, with a new generation process starting as soon as the entanglement resource is depleted – whereas reactive strategies aim at on-the-fly distribution of entanglement, with a new generation process starting on demand, when needed.

The choice between the two different strategies has a large impact on the network design, radically influencing the quantum network functinalities, somehow similarly to the choice between connection-oriented or connectionless services for classical networks.

Indeed, EPR pairs enable half-duplex unicast channels between pairs of nodes, regardless of their relative positions within the network topology. From this perspective, entanglement seems more reminiscent of connection-oriented circuit-switching rather than connection-less packet switching. Moreover, entanglement requires tight synchronization and signaling, unlikely satisfied by the best-effort nature of packet-switched networks. From the above, whether we should follow a completely packet-switching philosophy – with an infrastructure with no global central management and based on best effort strategy – or should we follow a more-oriented circuit switching philosophy – with an infrastructure similar to the telephone network that is based on central nodes in charge of network optimization and management – is a fundamental *philosophical* decision with system-wide cascade effects [1].

III. Opportunity

To face with the described challenges, a key role is played by standardization efforts of the different network functionalities, by abstracting from the particulars of the underlying technologies.

In this scenario, the capability of "moving" qubits out of the quantum nodes through the network for distributing entanglement [4] is mandatory. Specifically, as shown in Fig. 2, a (standard) transducer is needed to convert a matter qubit – that is, a qubit for information processing/storing within a computing device – in a flying qubit – which creates entanglement among remote nodes of the network – with a global consensus towards the use of photons as entanglement carriers [5]. Nowadays, there exist multiple technologies for

The authors are with the Quantum Internet research group, University of Naples Federico II (Web: www.quantuminternet.it). Emails: {angelasara.cacciapuoti, jessica.illiano, marcello.caleffi}@unina.it

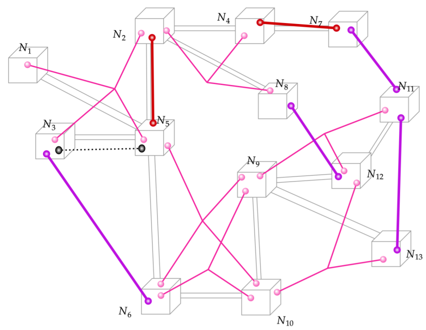

(a) Pictorial representation of the entanglement-based connectivity at a certain time instant.

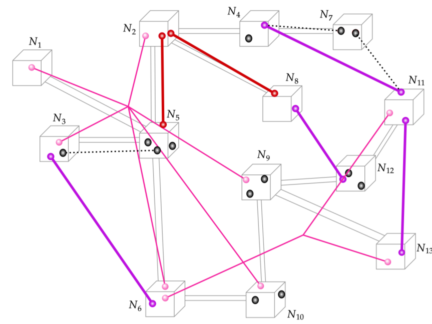

(b) Example of the network dynamics induced by EPR distillation at node N_4, entanglement swapping at nodes N_7, N_9, N_{10}, N_{12}, and entanglement merging at node N_5.

Fig. 1: Representation [1] of the dynamic changes in the netwrok connectivity enabled by the entanglement. By comparing Figure 1a and Figure 1b, it is evident that the number, the characteristics and the node identities of the virtual links are notably different, as a consequence of some LOCC operations, such as entanglement swapping, merging and distillation.

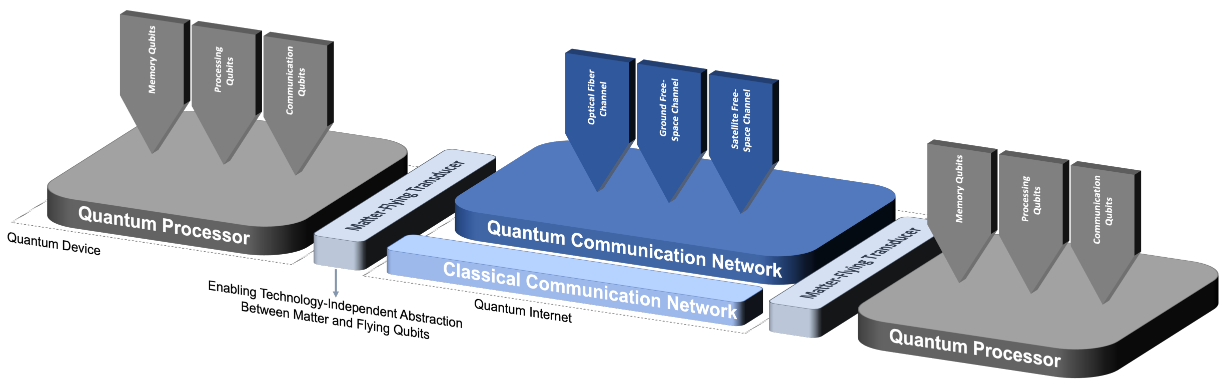

Fig. 2: Matter-Flying Interface for entanglement distribution within the quantum nodes of the network.

realizing a matter qubit, ranging from superconducting to ion-traps. And each technology is characterized by different pros and cons. As a consequence, a matter-flying interface would also face with this technology diversity. Besides, from a communication engineering perspective, the interface should be compatible also with the peculiarities of the physical channels the flying qubits propagate through. In fact, there exist different physical channels for transmitting flying qubits, ranging from free-space optical channels (either ground or satellite free-space) to optical fibers. Finally, the matter-flying interface should enable a technology independent abstraction in order to decouple underneath quantum hardware from upper software layers and to accelerate the development of industrial quantum ecosystems [1].

IV. Assessment

An efficient standardized transducer would constitute success towards the deployment of an entanglement-based network. Indeed, currently, there are several international projects and standardization efforts (e.g., in ITU, IETF, IEEE, GSMA, ETSI) which aim at defining architectures, interfaces and protocols for quantum networks. These diverse efforts should work in synergy to achieve the common and highly ambitious goal of an operating Quantum Internet.

References

[1] J. Illiano, M. Caleffi, A. Manzalini, and A. S. Cacciapuoti, "Quantum internet protocol stack: A comprehensive survey," *Computer Networks*, p. 109092, 2022.

[2] W. Dür, R. Lamprecht, and S. Heusler, "Towards a quantum internet," *European Journal of Physics*, vol. 38, no. 4, p. 043001, 2017.

[3] A. S. Cacciapuoti, M. Caleffi, R. Van Meter *et al.*, "When entanglement meets classical communications: Quantum teleportation for the quantum internet," *IEEE TCOM*, vol. 68, no. 6, pp. 3808–3833, 2020, invited paper.

[4] A. S. Cacciapuoti, M. Caleffi, F. Tafuri *et al.*, "Quantum internet: Networking challenges in distributed quantum computing," *IEEE Network*, vol. 34, no. 1, pp. 137–143, 2020.

[5] W. Kozlowski, S. Wehner, R. V. Meter, B. Rijsman, A. S. Cacciapuoti, M. Caleffi, and S. Nagayama, "Architectural Principles for a Quantum Internet," RFC 9340, Mar. 2023. [Online]. Available: https://www.rfc-editor.org/info/rfc9340

[6] C. Wang *et al.*, "Application Scenarios for the Quantum Internet," Internet Engineering Task Force, Internet-Draft draft-irtf-qirg-quantum-internet-use-cases-15, Mar. 2023, work in Progress.

[7] A. S. Cacciapuoti, J. Illiano, S. Koudia, K. Simonov, and M. Caleffi, "The quantum internet: Enhancing classical internet services one qubit at a time," *IEEE Network*, vol. 36, no. 5, pp. 6–12, 2022.

[8] J. Miguel-Ramiro *et al.*, "Genuine quantum networks with superposed tasks and addressing," *npj Quantum Information*, vol. 7, p. 135, 09 2021.

The Quantum Internet: Quest for a Paradigm Shift

Marcello Caleffi, *Senior Member, IEEE*, Jessica Illiano, Angela Sara Cacciapuoti, *Senior Member, IEEE*

I. Topic

Models: hybrid quantum and classical network design

II. Challenge

The design of the Quantum Internet [1]–[8] quests for a paradigm shift to properly harness the peculiarities of quantum entanglement and quantum information.

And indeed this paradigm shift cannot be limited to simply replace a few classical layers with some equivalent quantum layers, due to the intrinsic dissimilarities between classical and quantum information, as summarized by Table 1 in [1].

To elaborate further on and as detailed in [1], differently from classical information, quantum information irreversibly degrades over time as a consequence of the decoherence process. Hence, quantum information is characterized by hard temporal constraints, which may not fit with the latency characterizing the current Internet. Furthermore, while classical information can be freely read, quantum information is irreversibly altered by any measurement, according to the quantum measurement postulate. Moreover, whenever a quantum state is unknown it cannot be duplicated due to the no-cloning theorem. Conversely, when it comes to generate and distribute entangled states among network nodes, there is no restriction in repeatedly preparing a specific known entangled state, even though tighter interactions among the entangled nodes are mandatory. In fact, the nodes need to agree in advance on the specific entangled state to be first generated and then distributed.

Besides, bits and qubits can be considered *singleton*, namely, they both are self-contained entities, which have a meaning per-se. Conversely, entanglement is a correlation between multiple qubits. Indeed, not only a single entangled qubit is useless, but more implications emerge. First, there must be a tight cooperation between the network nodes – nodes that must be aware of each other identities – storing the entangled qubits for being able to exploit the quantum correlation provided by entanglement. Furthermore, any processing of a single entangled qubit has an instantaneous effect on the global entangled state, with possible changes affecting the remaining entangled qubits as well, regardless of the distances among the entangled nodes. Accordingly, entanglement exhibits a *non-local* scope. Conversely, both classical and quantum information – when flowing through the network for reaching the destination – exhibit local scope: any node can independently operate on it (as instance, to implement some error correction mechanisms) and the changes remain local.

The authors are with the Quantum Internet research group, University of Naples Federico II (Web: www.quantuminternet.it). Emails: {marcello.caleffi, jessica.illiano, angelasara.cacciapuoti}@unina.it

	Bit	Qubit	Entanglement
Temporal Constraints	**no**: can be stored indefinitely	**yes**: irreversibly degrades over time as a consequence of the decoherence process	
Duplication Constraints	**no**	**yes**: due to the no-cloning theorem	**no**: entangled states exploited in the network are in a known state, so they can be prepared repeatedly
Singleton	**yes**: self-contained entities		**no**: a single entangled qubit is useless in the network without the awareness of the remaining entangled qubits
Scope	**local**: any processing affects only the information available locally at the node		**non-local**: any processing of a single entangled qubit has an instantaneous effect on the remaining entangled qubits
State	**nearly stateless**: the node storing the bit does not need to retain any additional information	**stateful**: the node storing the qubit needs to retain at least temporal information	**profoundly stateful**: the node storing the entangled qubit needs to retain temporal information and the identities of the entangled nodes
Value	**local** and **pre-determined**: the encoded information is valuable only for the destination and not for the intermediate nodes		**global** and **dynamic**: the entangled state represents a valuable resource for any set of nodes sharing it
Order of Operations & Flow Direction	**yes**, with a strict ordering: source, intermediate nodes, destination	**flexible** the order: among the communication channels traversed by a quantum information carrier, can be indefinite	**flexible**: the swapping operation can happen simultaneously or without any particular order
Classes	**no**: there exist no classes of bits or qubits		**yes:** with a complex classification

Fig. 1. Summary of the main differences arising with quantum bits and quantum entanglement with respect to classical bits [1].

We underline that, when it comes to the design of the network functionalities, the difference between local and non-local scope is pivotal. With local scope, there is at any time a single network entity to whom the responsibility for the successful delivery of the information is delegated. Differently, non-local scope requires a tight coordination between multiple remotely-located peer entities. These peer entities may even compete among each others, as instance when multiple nodes simultaneously wish to use the same entanglement resource.

Another key dissimilarity is stateful vs stateless. Indeed, in the classical Internet, bits are usually transmitted in batch under the form of packets. Although some network functionalities acting on packets – with routing being a notable example – might require some sort of state information, bit *per-se* is stateless, i.e., the node storing the bit does not need to retain any additional information or detail for being able to operate on it. Conversely, the temporal constraints on qubits require some form of state information to be generated and distributed among the network entities. As instance, some temporal information regarding the quantum state residual coherence time must be available at the node for properly operating on it. Furthermore, the non-local scope characterizing entanglement requires additional state information – including at the very least the identities of the entangled nodes – to be properly shared through the network. Hence, while bits are nearly stateless, qubits and entangled qubits are definitely stateful.

Furthermore, it must be noted that entanglement constitutes

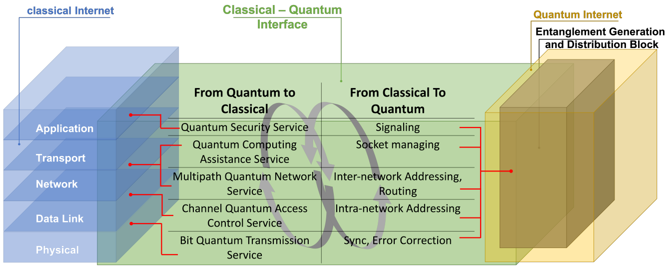

Fig. 2. Pictorial representation of the interplay between classical Internet and Quantum Internet [8].

an highly heterogeneous resource: there exist different classes of multipartite entangled states, which exhibit different properties and enable different network functionalities. And, such an heterogeneity has not been completely explored so far.

In conclusion, unconventional phenomena characterizing quantum mechanics completely twist the fundamental assumptions of several classical internet layers. And the intrinsic dissimilarities between qubits, entanglement and classical information affect the whole protocol stack, by even preventing the possibility of a one-to-one mapping between the classical network functionalities at a certain layer and the quantum ones within the designed protocol stack.

III. Opportunity

The Quantum Internet is unlikely to be functionally autonomous and independent of the classical Internet. Moreover, several functionalities – such as neighbor discovery, path discovery and forwarding – are spread among several layers as a consequence of the unique features of quantum entanglement described above. As a consequence, quantum communication protocols entail a dense cross-layer interdependence, which goes beyond the exchange of services between adjacent layers. Hence, the modeling given by a protocol stack – if possible – should be enriched by a system capable of implementing this wider cross-layer interaction.

Furthermore, the simplification given by the separation of concern on which the classical Internet is built on, and which groups functionalities into layers with adjacent-only layer interactions, seems unfeasible when it comes to quantum networks.

A possible solution would be to implement cross-layer interactions through classical signaling routed within the classical Internet, which would act as a unified interface to each layer of the quantum protocol stack. Clearly, with this solution further issues arise: should we exploit and adapt existing classical functionalities to implement quantum cross-layering, or do we need to design these functionalities from scratch [1], [8]?

Another solution – complimentary to the first one and mandatory whether cross-layer interactions should require exchange of quantum information – is to explicitly embed, within the same Quantum Internet protocol stack, cross-layer interactions among the quantum layers though the design of a classical-quantum interface, as represented in Figure 2.
Remarkably, the role of the aforementioned classical–quantum interface is not limited to enable cross-layer interactions within the Quantum Internet protocol stack. Specifically, not only the classical Internet offers services to the Quantum Internet, such as classical signaling. But, the Quantum Internet exhibits the potential of supporting and even enhancing classical internet functionalities as represented in Figure 2 from [8]. Concrete examples of this are analyzed in [8].

In this context, it is evident that a classical–quantum interface is needed to allow the a bidirectional interplay between the classical Internet and the Quantum Internet. As a matter of fact, the interplay between classical Internet and Quantum Internet cannot be limited to a single classical–quantum interface between a classical layer offering (or requiring) some specific service to a quantum counterpart layer. But it rather requires several interactions – likely differing in which part (quantum or classical) behaves as communication service provider – potentially involving different layers of the classical Internet protocol stack. As a consequence this interface should be a unified interface.

IV. Assessment

The Quantum Internet design is an ambitious long-term goal, with the lack of univocal metrics allowing a fair and quantitative comparison among different proposals. Certainly, the Quantum Internet design requires competences that span from the field of quantum physics, through computer science to communication and networking engineering. We firmly believe it will only emerge from the synergistic collaboration of researchers and companies with a multi-disciplinary approach.

References

[1] J. Illiano, M. Caleffi, A. Manzalini, and A. S. Cacciapuoti, "Quantum internet protocol stack: A comprehensive survey," *Computer Networks*, p. 109092, 2022.

[2] A. S. Cacciapuoti, M. Caleffi, F. Tafuri *et al.*, "Quantum internet: Networking challenges in distributed quantum computing," *IEEE Network*, vol. 34, no. 1, pp. 137–143, 2020.

[3] S. Wehner, D. Elkouss, and R. Hanson, "Quantum Internet: a Vision for the Road Ahead," *Science*, vol. 362, no. 6412, 2018.

[4] R. Van Meter, R. Satoh, N. Benchasattabuse, T. Matsuo, M. Hajdušek, T. Satoh, S. Nagayama, and S. Suzuki, "A quantum internet architecture," *arXiv preprint arXiv:2112.07092*, 2021.

[5] W. Kozlowski, S. Wehner, R. V. Meter, B. Rijsman, A. S. Cacciapuoti, M. Caleffi, and S. Nagayama, "Architectural Principles for a Quantum Internet," RFC 9340, Mar. 2023.

[6] J. Miguel-Ramiro, A. Pirker, and W. Dür, "Optimized quantum networks," *Quantum*, vol. 7, p. 919, 2023.

[7] W. Dür, R. Lamprecht, and S. Heusler, "Towards a quantum internet," *European Journal of Physics*, vol. 38, no. 4, p. 043001, 2017.

[8] A. S. Cacciapuoti, J. Illiano, S. Koudia, K. Simonov, and M. Caleffi, "The quantum internet: Enhancing classical internet services one qubit at a time," *IEEE Network*, vol. 36, no. 5, pp. 6–12, 2022.

The Need for End-to-End Application Performance Benchmarking of Quantum Hardware

Daan Camps[1], Siva Darbha, Katherine Klymko, Ermal Rrapaj, Nicholas J. Wright

National Energy Research Scientific Computing Center (NERSC), LBNL

Topics: Applications, Algorithms

Introduction and Background. The ongoing global effort to develop scalable quantum computing hardware platforms aims to overcome some of the grand computational challenges of the 21st century as the rate of improvement in classical compute infrastructure is expected to slow down [1]. Specifically, applications that study phenomena that are quantum mechanical in nature, e.g. materials science, chemistry, and low-energy nuclear physics, are expected to benefit from quantum computing hardware. As such, it is crucial for the mission of the Advanced Scientific Computing Research (ASCR) program of the US Departmenet of Energy (DOE) Office of Science to (i) better understand the performance of quantum hardware platforms through application-inspired benchmarks and curated libraries of instances, and (ii) explore new DOE science applications that may be accelerated by quantum computers.

Challenge. Future quantum systems will likely serve as accelerators in HPC data centers [2, 3], running hybrid quantum-classical applications on a tightly-coupled, distributed, hybrid quantum-classical (DHQC) system. Hybrid quantum-classical computational kernels have been proposed [4], such as variational algorithms and subspace methods [5]. However, at present, very few end-to-end science applications exist that truly benefit from a DHQC platform, primarily because of the limited capabilities of quantum hardware today. At the same time, application-level benchmarks have been proposed [6], but they remain limited as they only test the core computational quantum kernel through a quantum circuit execution, instead of the full end-to-end application. As quantum computing hardware improves, a primary challenge will be to develop application-level benchmarks that probe the aggregate performance of hybrid quantum-classical infrastructure, a critical step in assessing realistic quantum advantages. A secondary challenge includes a systematic exploration of the problem space relevant to the Office of Science to identify new kinds of science applications that are currently intractable to classical HPC but can be enabled by future quantum capabilities. Notably, this includes revisiting existing algorithms that can be accelerated on DHQC computers [5, 7].

Opportunity. We describe how the two challenges above may be addressed over the next years:
(i) End-to-End Application Performance Benchmarking—The following milestones will have to be met in order to develop a end-to-end applications and benchmark their performance on a DHQC system. Firstly, a quantum computing platform will have to be co-located and integrated with Office of Science computing facilities. Secondly, it is of the highest importance that domain scientists, QIS experts, and HPC engineers work closely together to identify the most promising existing application(s) that can be ported this platform. We believe that each of these three groups are critical for success: domain scientists can identify the important science problems and state-of-the-art classical codes, QIS experts possess a deep knowledge of the strengths and limitations of quantum capabilities, and HPC engineers hold extensive technical expertise in HPC architecture and performance metrics and benchmarking [8]. It is only by such a deep collaborative effort that we can expect an end-to-end application to be realized on an early DHQC system at an Office of Science compute facility.
(ii) Exploration of new science applications—A prominent example of science applications that are

[1]corresponding author: dcamps@lbl.gov

accelerated by quantum computers is quantum dynamics. Where classically simulating quantum dynamics requires exponentially increasing memory and time with problem size, quantum computers are likely perform these simulations with moderate resources for many relevant problems. We identify an opportunity to assess which existing applications may be reformulated in terms of a quantum dynamics and which novel applications can be explored through quantum dynamics. Two examples of the former include the study of entanglement and dynamical phase transitions during collective neutrino flavor oscillations in supernovae and neutron star binary mergers [9, 10, 11] and the computation of eigenenergies of chemical systems using a subspace of real-time expansion states [5]. Other such applications likely exist and can be discovered through systematic exploration.

Assessment. A demonstration of end-to-end science applications executed on a DHQC system located in an Office of Science compute facility is an important measure of success for the first challenge. Additional metrics include a quantitative comparison of the DHQC applications with classical HPC in terms of scaling performance, time-to-solution, and energy usage. This would constitute an important milestone to assess the utlity of quantum hardware. Furthermore, this effort has the potential to inform further investment decisions about the development and procurement of quantum hardware. The discovery of new science applications may be informed by the end-to-end performance benchmark developed in the first challenge. This challenge will be considered successful if new science applications, e.g., enabled through simulations of quantum dynamics, are included in the Office of Science portfolio.

Timeliness or maturity. Quantum hardware platforms are rapidly approaching a level of maturity that will allow them to shift from the lab space to a data center environment. At the same time, the classical HPC space is exploring beyond-Moore technology options. It is crucial to address the two challenges we identified in this position paper in the next few years in order to assess the viability of quantum technologies as a scalable computational resource.

References

[1] T. N. Theis and H.-S. P. Wong, "The end of Moore's law: A new beginning for information technology," *Computing in Science Engineering*, vol. 19, no. 2, pp. 41–50, 2017.

[2] IBM Quantum, "Quantum-centric supercomputing: The next wave of computing." `https://research.ibm.com/blog/next-wave-quantum-centric-supercomputing`, 2022. Accessed: 2023-04-24.

[3] K. A. Britt, F. A. Mohiyaddin, and T. S. Humble, "Quantum accelerators for high-performance computing systems," in *2017 IEEE International Conference on Rebooting Computing (ICRC)*, pp. 1–7, 2017.

[4] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, "Hybrid quantum-classical algorithms and quantum error mitigation," *Journal of the Physical Society of Japan*, vol. 90, no. 3, p. 032001, 2021.

[5] K. Klymko, C. Mejuto-Zaera, S. J. Cotton, F. Wudarski, M. Urbanek, D. Hait, M. Head-Gordon, K. B. Whaley, J. Moussa, N. Wiebe, W. A. de Jong, and N. M. Tubman, "Real-Time Evolution for Ultracompact Hamiltonian Eigenstates on Quantum Hardware," *PRX Quantum*, vol. 3, p. 020323, May 2022.

[6] T. Tomesh, P. Gokhale, V. Omole, G. Ravi, K. N. Smith, J. Viszlai, X. Wu, N. Hardavellas, M. R. Martonosi, and F. T. Chong, "Supermarq: A scalable quantum benchmark suite," in *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, (Los Alamitos, CA, USA), pp. 587–603, IEEE Computer Society, apr 2022.

[7] D. Camps, R. Van Beeumen, and C. Yang, "Quantum Fourier transform revisited," *Numerical Linear Algebra with Applications*, vol. 28, no. 1, p. e2331, 2021.

[8] B. Austin, C. Daley, D. Doerfler, J. Deslippe, B. Cook, B. Friesen, T. Kurth, C. Yang, and N. J. Wright, "A metric for evaluating supercomputer performance in the era of extreme heterogeneity," in *2018 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pp. 63–71, 2018.

[9] E. Rrapaj, "Exact solution of multiangle quantum many-body collective neutrino-flavor oscillations," *Phys. Rev. C*, vol. 101, p. 065805, Jun 2020.

[10] A. Roggero, E. Rrapaj, and Z. Xiong, "Entanglement and correlations in fast collective neutrino flavor oscillations," *Phys. Rev. D*, vol. 106, p. 043022, Aug 2022.

[11] A. V. Patwardhan, M. J. Cervia, E. Rrapaj, P. Siwach, and A. B. Balantekin, *Many-Body Collective Neutrino Oscillations: Recent Developments*, pp. 1–16. Singapore: Springer Nature Singapore, 2020.

Post-Quantum Cyberinfrastructure Security Readiness: Risks, Measures and Prospects

Phuong Cao[1,2] – pcao3@illinois.edu
Bach Hoang[3] – bachh2@illinois.edu
Santiago Núñez-Corrales[1,2] – nunezco2@illinois.edu
[1]National Center for Supercomputing Applications, UIUC
[2]Illinois Quantum Information Science and Technology Center, UIUC
[3]Department of Mathematics, UIUC

The problem of correctly implementing quantum resistant cryptographic network protocols is critically important to drive the adoption of quantum computers to the masses. It is urgent because practical quantum computers are on the horizon: 1,000-QUBIT quantum computer [1] will soon be offered by major industry players such as IBM in 2023, while some nations also claimed to break RSA encryption already [3]. The main challenges of implementing post quantum cryptography are: 1) complexity of implementing and verifying new cryptographic protocols, 2) wide-spectrum of client's cryptographic protocols such as IoT devices, and 3) significant performance overhead when deploying on Internet-scale networked computers. Nevertheless, how existing cyberinfrastructure will support post-quantum cryptography is largely unknown.

This paper proposes a testbed of novel networked telescopes that will be deployed at the nation's backbone (Figure 1). The networked telescopes, deployed at different vantage points, will continuously measure PQC adoption, characterize the performance overhead, and provide a real-time feedback loop to NIST in order to improve and fix potential security bugs of PQC algorithms in development. By bringing a diverse team of a cybersecurity expert, a mathematician, and an information theorist, we will initiate intellectual discussions at the workshop. This distributed network of telescope is our opportunity to stay ahead of attackers.

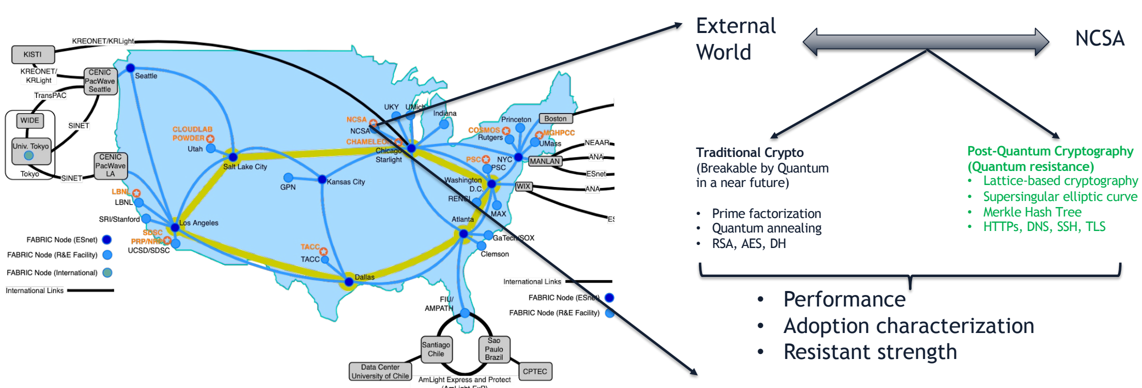

Figure 1: Networked telescope testbed to characterize deployment progress and performance overhead of post quantum cryptography algorithms.

Methods. The networked telecope testbed in Figure 1 will advances our understanding of quantum-resistant cryptographic algorithms are being deployed in the wild, e.g., TLS Post-Quantum Confidentiality key exchange algorithm in TLS (CECPQ2). Our approach is to build an array of a globally distributed network telescopes, each telescope is placed at a network border router to tap into incoming/outgoing network traffic. For example, the National Center for Supercomputing Applications (NCSA) at the University of Illinois has already been collecting network connection metadata, e.g., the packet headers, hand-shake algorithms, and cryptographic suite in encrypted protocols such as TLS, SSH, and encrypted DNS. The networked telescope will continuously measure, scan and remedy weaknesses in real world quantum implementations, as well as exchange of threat intelligence. We plan to deploy our master telescope at NCSA, a choke point of scientific traffic that would provide 360-degree, 24/7, orthogonal view of scientific network traffic. We will capture the network traffic through 400Gbps network border link. Using NIST's recommended algorithms in quantum cryptography such as CRYSTALS-Kyber for encryption, FALCON and SPINCS+ for digital signature, and KEMTALS for key exchange, we will test the usage of these algorithms in modern network protocols such as HTTPS (TLS 1.3), SSH, DNSSEC, and QUIC. This master telescope, when being connected with others in the future, will provide real-time Internet-wide scans that characterize the upcoming deployment of NIST's recommended encryption and digital signature algorithms for Post Quantum Cryptography.

Broader impact of our approach. The opportunity is to identify critical weaknesses or vulnerable implementations of quantum resistant cryptographic protocol in real time before the attackers can exploit them. Whether to fully adopt it is a controversial topic because of the added computational complexity, performance overhead, and unknown security issues. The success metric is the percentage of devices including IoT devices that correctly implement and support quantum resistant cryptography 100% of the time. Our solution is to build a globally distributed network telescope to continuously measure , scan and remedy weaknesses in real world quantum implementations.

Putting our approach in perspective. Recently, researchers have identified two Post-Quantum Algorithms that have the nest performance are Dilithium and Falcon, while Falcon seems to be more suitable for the web. One of the existing prooblem about Post-Quantum crytography when integrating with the network server is the signing, when slightly slower signing can have significant impact on the whole server. The concentration is now moved to signature and key size, as the optimization and hardware accleration improve signing performance. Some developments have been made by combining some Post-Quantum Cryptography algorithms to improve the handshake speed and leveraging ICA suppression to avoid round-trips. The future for Post-Quantum Crytography intergration to network server is to test the performance of PQ authenticated VPNs and UDP-based tunnles like QUIC and DTLS. Some experiments should be conducted to quantify the total impact of PQ algorithm under realistic conditions that include lossy networks. In addition, investigation of hybrid certificates' performance should be carried out by studying the message recovery capabilities offered by schemes such as Falcon. [2]

References

[1] An ibm quantum computer will soon pass the 1,000-qubit mark - ieee spectrum. https://spectrum.ieee.org/ibm-condor. (Accessed on 04/30/2023).

[2] George Tasopoulos, Charis Dimopoulos, Apostolos P. Fournaris, Raymond K. Zhao, Amin Sakzad, and Ron Steinfeld. Energy consumption evaluation of post-quantum tls 1.3 for resource-constrained embedded devices. Cryptology ePrint Archive, Paper 2023/506, 2023. https://eprint.iacr.org/2023/506.

[3] Bao Yan, Ziqi Tan, Shijie Wei, Haocong Jiang, Weilong Wang, Hong Wang, Lan Luo, Qianheng Duan, Yiting Liu, Wenhao Shi, Yangyang Fei, Xiangdong Meng, Yu Han, Zheng Shan, Jiachen Chen, Xuhao Zhu, Chuanyu Zhang, Feitong Jin, Hekang Li, Chao Song, Zhen Wang, Zhi Ma, H. Wang, and Gui-Lu Long. Factoring integers with sublinear resources on a superconducting quantum processor, 2022.

Title: Co-Design Quantum Memories for Quantum Networks

Authors: Matthew Cha[1], Alex Haid, Mike Jaris, Pejman Jouzdani, Mark Kostuk, Nils Paz, David Schissel, Dan Sioui
General Atomics, 3550 General Atomics Ct., San Diego, CA 92121
[1] Corresponding author: chamatthew@fusion.gat.com

Topic: Quantum networking; Codesign and Integration

Challenge: Quantum networks enable long-distance quantum communication. The nodes of a quantum network vary, and range from simple single qubit systems for secure communication or powerful quantum computing processors for distributed quantum computation. In each case, robust multi-qubit quantum memories are necessary for enabling scalable quantum networks where they may function as quantum repeaters [1], for quantum error detection and quantum error correction [2, 3], and as random-access quantum memory [4]. An integrated framework for the co-design of quantum memories for quantum networks will enable a rapid advance of quantum materials and quantum control essential in the transition from the current experimental quantum networks to scalable quantum networks with robust functionality and application.

One key challenge is to preserve the static properties of a quantum memory within dynamic quantum network protocols. For example, Si-V color center qubits have demonstrated quantum network primitives such as entanglement distillation and entanglement swapping with error detection [2]. A key challenge in this experiment was to preserve the static quantum memory properties of Si-V color center while photonic qubits are coupled into and out of the material systems for quantum communication. New sources of decoherence arise in the form of stochastic ionization, spin crosstalk, and leakage, and in some cases, improved quantum characterization and control could preserve robustness properties.

A second challenge is characterizing the necessary requirements of a quantum memory to enable quantum network functionality and applications, and improving the selection criteria for quantum memory to optimize the performance on specific quantum network protocols. The above experiment, and similar ones [1, 3], show the promise of quantum memories for scalable quantum repeaters within quantum networks. However, characterizing the transition from experimental results to full functionality and applications will require rapid development in tools for quantum network characterization, quantum network protocol requirements, and simulation models incorporating quantum memory parameterization.

Opportunity: The integration and performance of quantum memories in quantum networks through quantum material co-design and implementation of novel quantum control is a priority research direction. In quantum materials, novel material synthesis and nanofabrication methods, such as ion implantation and chemical vapor deposition, have led to record quantum memory lifetimes in materials such as defects in SiC and color centers in diamond [5]. The theoretical discovery of dynamically driven topological phases has opened new avenues for qubit modalities as quantum memories, such as in trapped-ion devices [6], and for improved quantum control and parametrization of quantum memories.

An opportunity exists for the co-design of material and theoretical developments in quantum memories for robust functionality and application to scalable quantum networks. Methods for verification and validation will need to be developed, such as metrics for link efficiency [3] to characterize the relationship of material parameterization to quantum network performance. Such an integrated framework would allow for focused material design and characterization of various quantum network applications and drive rich collaborations between research groups in quantum materials and quantum networks.

Assessment: Quantum network testbeds have recently been formed across the United States including the Q-NEXT in Chicago, QUANT-NET in northern CA, and NSF Center for Quantum Networking Boston Testbed. This enables the ability to rapidly test and prototype potential solutions for quantum memories deployed in experimental quantum networks using various quantum network protocols and interacting with multiple qubit modalities. Distinct quantum memory platforms will arise to give optimal performance for each functionality such as quantum repeaters, quantum error detection, and QRAM.

Timeliness and maturity: Breakthroughs in quantum memory properties and lifetimes have been recently achieved in a variety of quantum materials [5, 6]. However, it is still unclear how these quantum memories will perform within a dynamic quantum network protocol, and further, what quantum network functionality and applications will be enabled. It is our position that the co-design of quantum memories for quantum network will enable the next major breakthrough in the functionality and applications of quantum networks with the fault tolerant quantum communication on the horizon.

Reference:

[1] Wehner, S., Elkouss, D. and Hanson, R., 2018. Quantum internet: A vision for the road ahead. Science, 362(6412).

[2] Stas, P.J., Huan, Y.Q., Machielse, B., Knall, E.N., Suleymanzade, A., Pingault, B., Sutula, M., Ding, S.W., Knaut, C.M., Assumpcao, D.R. and Wei, Y.C., 2022. Robust multi-qubit quantum network node with integrated error detection. Science, 378(6619).

[3] Bradley, C.E., de Bone, S.W., Möller, P.F.W., Baier, S., Degen, M.J., Loenen, S.J.H., Bartling, H.P., Markham, M., Twitchen, D.J., Hanson, R. and Elkouss, D., 2022. Robust quantum-network memory based on spin qubits in isotopically engineered diamond. npj Quantum Information, 8(1).

[4] Jiang, N., Pu, Y.F., Chang, W., Li, C., Zhang, S. and Duan, L.M., 2019. Experimental realization of 105-qubit random access quantum memory. npj Quantum Information, 5(1).

[5] de Leon, N.P., Itoh, K.M., Kim, D., Mehta, K.K., Northup, T.E., Paik, H., Palmer, B.S., Samarth, N., Sangtawesin, S. and Steuerman, D.W., 2021. Materials challenges and opportunities for quantum computing hardware. Science, 372(6539).

[6] Dumitrescu, P.T., Bohnet, J.G., Gaebler, J.P., Hankin, A., Hayes, D., Kumar, A., Neyenhuis, B., Vasseur, R. and Potter, A.C., 2022. Dynamical topological phase realized in a trapped-ion quantum simulator. Nature, 607(7919).

High-performance non-Gaussian resource state generation for quantum repeaters and distributed quantum computing

J. C. Chapman,[1,*] R. C. Pooser,[1] A. M. Marino,[1] and N. A. Peters[1]

[1] *Quantum Information Science Section, Oak Ridge National Laboratory, Oak Ridge, TN*

[*] *chapmanjc@ornl.gov*

Topics. (primary) models and (secondary) error correction and mitigation.

Challenge. Long distance quantum networking requires methods for detecting and correcting errors resulting from transmission loss and operational errors to enable high-performance. These errors can theoretically be corrected using quantum repeaters [1]—a technology proposed specifically to solve this problem. Quantum repeater proposals vary widely in how they achieve this technological goal, but almost all use optical photons as the carriers of quantum information to achieve efficiency and speed. Moreover, an all-optical quantum repeater [2] could avoid the need for stationary qubits (e.g., in a matter-based quantum memory), which would improve speed and efficiency. All-optical quantum information processing generally has two models of design: (1) create simple few-mode states and combine them with more difficult multi-mode operations; (2) create more complex multi-mode states and manipulate them with simpler few-mode operations. In either case, the difficulty arises from the inefficient non-linear multi-mode interactions that can only be achieved with the help of a physical interaction medium, e.g., a nonlinear crystal. Engineering strong interactions between photons, whether to generate large entangled states or to implement efficient multi-photon operations, has been a challenge in realizing all-optical quantum repeaters and quantum computing.

Opportunity. Since non-linear multi-photon interactions are challenging, limiting their use is advantageous. Thus, the above design model (2) is more likely to provide a lower implementation complexity and potentially lower cost by focusing the increased resources on the state generation at centralized locations. Of the complex resource state proposals [2–5], the ones using hybrid continuous/discrete variable (CV/DV) non-Gaussian states [4,5]—in particular, those based on Gottesman-Kitaev-Preskill (GKP) qubits [6]—offer several attractive advantages over other repeater proposals.

Non-Gaussian states' hybrid nature can leverage the advantages of each of the previously mentioned approaches while minimizing their disadvantages—as pointed out in a previous ASCR workshop report [7]—to lead to new capabilities for all-optical quantum repeaters and quantum computers. In particular, CV systems offer unconditional operation, high detection efficiency through homodyne detection, and more practical interfacing with a number of quantum systems; however, they suffer from sensitivity to losses and limited fidelities. On the other hand, DV systems can achieve high fidelity operation and enable entanglement distillation after losses; however, their operation is probabilistic. The ability to use homodyne detectors, in particular, enables lower C-SWAP (cost, size, weight, and power) compared to most single-photon detection systems, high efficiency operation at room temperature, and filtering that can enable high-performance coexistence with conventional networking [8, 9]. Additionally, the generation of GKP states requires resource ancilla with large squeezing that can be efficiently and deterministically generated with room-temperature spontaneous parametric downconversion [10].

GKP encoding, in particular, provides a natural way to deal with loss errors better than some other proposals specifically designed to correct loss errors [11]; this is due in part to the multiple discrete phase-space features of the state that comprise a sort of grid and make this a hybrid CV/DV state. These error-handling properties make GKP encoding a good candidate for quantum repeaters as losses are a dominant noise source in optical networks. As an example, GKP states produced in superconductors were recently used for error correction beyond the break-even point [12]. By leveraging homodyne detection and squeezing, proposals using hybrid CV/DV non-Gaussian resource states are well-suited to develop quantum repeaters for commercial deployment with lower complexity and cost than alternative proposals.

The optical generation of non-Gaussian resource states, e.g., GKP qubits, remains a challenge. Recent proposals to create such resource states starting with smaller available optical resources, improving them, and combining them in several new ways to create the desired resource state [5, 13–17] provide an opportunity to address this challenge. Undoubtedly, increased experimentation in this area will help to refine designs for generating non-Gaussian resource states and bring the high-quality generation of those states closer to reality.

Assessment. For quantum repeaters, the key metric of success is whether the transmission rate and transmitted-state quality through the repeater network are better than directly transmitting the input state over the same distance without repeaters. In practice, the resources required at each repeater node and the required node spacing are also important to consider. More specific to developing hybrid CV/DV non-Gaussian resource states, key metrics include the negativity level of the state's Wigner function, the squeezing level of the states used to generate them and their separation in phase space, or the number of negative dips in the Wigner function. These figures of merit provide: indications of the departure from a classical state, enable error correction encodings, and potentially higher error tolerance. Tolerable thresholds for various metrics are usually specific to a

given repeater proposal. In general, the key metrics for a particular resource state depend on the quantum repeater protocol.
Timeliness or maturity. Multiple recent proposals to generate non-Gaussian resources are promising in their practicality. For example, cat state generation, which produces a superposition of coherent states, via repeat-until-success photon subtraction [13] can provide a practically deterministic non-Gaussian resource for further processing via simple, fully deterministic Gaussian operations. New "breeding" protocols [14] then process these cat states further to produce error-correctable GKP states. Successful demonstration of these protocols would provide the necessary ingredients for an error-corrected repeater. Recent technological developments, especially in relation to squeezing sources [10] and distributed joint homodyne detection [9] bring us closer to implementing non-Gaussian resource state generation to enable quantum repeaters. These initial experiments help identify the remaining gaps and help inform how to bridge them.

Given that GKP states can serve as the basis for an error-corrected repeater, while also expressing a hybrid CV-DV qubit, synthesis of these states would provide a valuable resource for distributed quantum computation. Because repeaters are used to distribute quantum entanglement amongst nodes on a network, the repeaters can themselves serve as the resource state for measurement-based quantum computing. If every node of the network contained a GKP state-based quantum repeater, then every node can be treated as a collection of error-corrected qubits entangled with a subset of the rest of the network. If the network is arranged in a simple cluster-state configuration, then the network with GKP-based repeaters becomes a distributed one-way quantum computer. Distributed quantum computing across nodes on a quantum network is the quintessence of what it means to build a "quantum internet," which promises to make a distributed quantum science infrastructure exponentially more powerful than a conventional network of its components.

References

1. H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, "Quantum repeaters: The role of imperfect local operations in quantum communication," Phys. Rev. Lett. **81**, 5932–5935 (1998).
2. K. Azuma, K. Tamaki, and H.-K. Lo, "All-photonic quantum repeaters," Nat. Commun. **6**, 6787 (2015).
3. F. Ewert, M. Bergmann, and P. van Loock, "Ultrafast long-distance quantum communication with static linear optics," Phys. Rev. Lett. **117**, 210501 (2016).
4. K. Fukui, R. N. Alexander, and P. van Loock, "All-optical long-distance quantum communication with Gottesman-Kitaev-Preskill qubits," Phys. Rev. Res. **3**, 033118 (2021).
5. F. Rozpedek, K. P. Seshadreesan, P. Polakos, L. Jiang, and S. Guha, "All-photonic multiplexed quantum repeaters based on concatenated bosonic and discrete-variable quantum codes," arXiv:2303.14923 (2023).
6. D. Gottesman, A. Kitaev, and J. Preskill, "Encoding a qubit in an oscillator," Phys. Rev. A **64**, 012310 (2001).
7. T. Ndousse-Fetter, N. Peters, W. Grice, P. Kumar, T. Chapuran, S. Guha, S. Hamilton, I. Monga, R. Newell, A. Nomerotski, D. Towsley, and B. Yoo, "Quantum networks for open science," arXiv:1910.11658 (2019).
8. B. Qi, W. Zhu, L. Qian, and H.-K. Lo, "Feasibility of quantum key distribution through a dense wavelength division multiplexing network," New J. Phys. **12**, 103042 (2010).
9. J. C. Chapman, A. Miloshevsky, H.-H. Lu, N. Rao, M. Alshowkan, and N. A. Peters, "Two-mode squeezing over deployed fiber coexisting with conventional communications," arXiv:2304.10053 (2023).
10. T. Kashiwazaki, T. Yamashima, K. Enbutsu, T. Kazama, A. Inoue, K. Fukui, M. Endo, T. Umeki, and A. Furusawa, "Over-8-dB squeezed light generation by a broadband waveguide optical parametric amplifier toward fault-tolerant ultra-fast quantum computers," arXiv:2301.12658 (2023).
11. V. V. Albert, K. Noh, K. Duivenvoorden, D. J. Young, R. T. Brierley, P. Reinhold, C. Vuillot, L. Li, C. Shen, S. M. Girvin, B. M. Terhal, and L. Jiang, "Performance and structure of single-mode bosonic codes," Phys. Rev. A **97**, 032346 (2018).
12. V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, "Real-time quantum error correction beyond break-even," Nature **616**, 50–55 (2023).
13. K. Marshall, R. Pooser, G. Siopsis, and C. Weedbrook, "Repeat-until-success cubic phase gate for universal continuous-variable quantum computation," Phys. Rev. A **91**, 032321 (2015).
14. D. J. Weigand and B. M. Terhal, "Generating grid states from Schrödinger-cat states without postselection," Phys. Rev. A **97**, 022341 (2018).
15. I. Tzitrin, J. E. Bourassa, N. C. Menicucci, and K. K. Sabapathy, "Progress towards practical qubit computation using approximate Gottesman-Kitaev-Preskill codes," Phys. Rev. A **101**, 032315 (2020).
16. K. Fukui, M. Endo, W. Asavanant, A. Sakaguchi, J.-i. Yoshikawa, and A. Furusawa, "Generating the Gottesman-Kitaev-Preskill qubit using a cross-Kerr interaction between squeezed light and Fock states in optics," Phys. Rev. A **105**, 022436 (2022).
17. R. Yanagimoto, R. Nehra, R. Hamerly, E. Ng, A. Marandi, and H. Mabuchi, "Quantum nondemolition measurements with optical parametric amplifiers for ultrafast universal quantum information processing," PRX Quantum **4**, 010333 (2023).

Title: Quantum Inspired Electric Grid Optimization with Uncertainties

Authors:
Yousu Chen, Pacific Northwest National Laboratory, yousu.chen@pnnl.gov
Ang Li, Pacific Northwest National Laboratory, ang.li@pnnl.gov
Meng Yue, Brookhaven National Laboratory, yuemeng@bnl.gov
Tong Geng, University of Rochester, tgeng@UR.Rochester.edu

Topic: Applications, models, and algorithms for quantum-inspired electric grid optimization

Challenge: Optimization of electric grid network is a challenging task because of its high nonlinearity, high dimensionality, and the involvement of multiple objectives covering various aspects of electric grid. The complexity of electric grid network optimization increases exponentially with the increasing penetration of renewables, increasing number of variables and constraints, which results in longer computational time. Meanwhile, the more dynamic system requires the quicker response to a disturbance or an event. Classical optimization methods face limitations in handling these complexities, leading to suboptimal solutions or cannot find a solution within a required time frame. For instance, it takes about 3 to 4 hours to find a solution of a security constrained unit commitment (SCUC) problem for a 45K-bus power system, with 200K constraints, 500K variables, and 50K binary variables. Solving this type of optimization problem with high dimensionality and non-convex constraints can be challenging for traditional techniques, even with the help from classical high-performance computing.

Opportunity: Quantum computing (QC) is an emerging computing technology that offers the potential to solve complex, large problems that are beyond the capabilities of traditional computers, without dramatically increasing the computational time. However, due to the current maturity of QC, it remains challenging to demonstrate the benefits for practical large-scale problems. For power system optimization, quantum-inspired algorithms, such as the adiabatic quantum optimization algorithm, or quantum approximated optimization algorithm (QAOA) would be promising in helping address the computational challenges. By leveraging quantum computing principles, these algorithms can explore the solution space more efficiently and find optimal solutions faster, potentially overcoming the limitations of classical methods. Therefore, quantum-inspired algorithms have a good potential to develop practical quantum computing solutions to large-scale power system optimization problems.

Assessment: Achieving success in this area would require the development and implementation of quantum-inspired algorithms that demonstrate superior performance and scalability compared to classical methods. Therefore, the assessment of these algorithms can be based on the following criteria:

(1) Solution quality: the solution quality of quantum-inspired algorithms should be better or at least as accurate as classical methods.
(2) Computational speed: The algorithm should be able to find a solution faster than classical methods, with an expectation of at least 10x speedup.
(3) Scalability: The algorithm should be scalable and able to handle large-scale systems without significant performance degradation.
(4) Robustness: The algorithm should be applicable to various operating conditions, including uncertainties and contingencies.

Timeliness or Maturity: With the rapid advancements in quantum computing and growing interest in quantum-inspired optimization, there is a unique opportunity to explore the potential of quantum-inspired algorithms in solving practical, large-scale power system optimization problems. The increasing complexity of power systems and the need for sustainable, reliable, and efficient operations, make it crucial to develop innovative optimization techniques that can address these challenges. By leveraging quantum-inspired algorithm-based approaches, we can make significant short-term impacts for large-scale problems and gain a better understanding of the bottleneck and opportunities for quantum computing applications in electric grid network, and most likely can be extended to other complex networks. Success in this area would contribute significantly to accelerate the renewable integration process by unleashing unprecedented computing power equipped with quantum technologies. This research direction also aligns well with the OE's objectives and missions and helps meet the clean energy goals of the administration.

References:

1. Arrazola, J.M., Delgado, A., Bardhan, B.R. and Lloyd, S., 2019. Quantum-inspired algorithms in practice. *arXiv preprint arXiv:1905.10415*.
2. Giani, A., Eldredge, Z. Quantum Computing Opportunities in Renewable Energy. SN COMPUT. SCI. 2, 393 (2021). https://doi.org/10.1007/s42979-021-00786-3
3. Hari P. Paudel, Madhava Syamlal, Scott E. Crawford, Yueh-Lin Lee, Roman A. Shugayev, Ping Lu, Paul R. Ohodnicki, Darren Mollot, and Yuhua Duan, Quantum Computing and Simulations for Energy Applications: Review and Perspective, ACS Engineering Au 2022 2 (3), 151-196, DOI: 10.1021/acsengineeringau.1c00033
4. H. Zheng, A. Khodaei, R. Eskandarpour, A. Paaso, "A Quantum Leap Is Coming: Ones, Zeros And Everything In Between," T&D World Magazine, Jan. 2021.
5. Ross, O.H.M., 2019. A review of quantum-inspired metaheuristics: Going from classical computers to real quantum computers. *IEEE Access*, *8*, pp.814-838.

Physics-Aware, Full-Stack Software Optimizations

Fred Chong
Seymour Goodman Professor, University of Chicago
Chief Scientist for Quantum Software, Infleqtion

I. INTRODUCTION

Despite sustained progress in quantum hardware, there is a substantial gap to utility-scale quantum computation. On one hand, there has been consistent improvement in gate fidelities over the past decade. For instance, since the advent of superconducting transmon qubits [1], two-qubit gate errors have been lowered by $\sim 0.77\times$ per year [2]. We see similar rates of progress in other qubit technologies including neutral atoms [3], [4] and trapped ions [5], [6]. While this progress is encouraging, hardware progress alone would require at least a decade to achieve societally-useful outcomes such as simulating molecules relevant to fertilizer production [7].

Quantum software can be a force multiplier that can significantly shorten the timeline for utility-scale results from quantum hardware. There are compelling parallels to classical computing: the world's top computing facilities bolster their state-of-the-art hardware capabilities, with significant investment in software tools such as CUDA [8], OpenMP [9], and SLURM [10]. Similarly, software tooling—especially for compilation—can enable users to extract better results from quantum hardware, both for near-term systems as well as for upcoming large-scale fault-tolerant computers.

In fact, we find even stronger motivation for optimized compilation in the quantum setting than in the classical setting. First, quantum resources are far more expensive than classical resources. Second, for foreseeable quantum computers, optimized compilation will be *necessary* to bring useful applications within the boundary of achievable computations. Lastly, applications brought within this boundary will exhibit exponential or high-degree polynomial quantum advantage [11], that can immediately justify high compilation costs. Thus, investment in deep compiler optimization can enable applications that are otherwise out of reach for current hardware at various scales.

In particular, several key research directions will help realize practical quantum advantage. Physics-aware, cross-layer optimizations will continue to yield important efficiencies to allow applications to make the most of quantum resources. Software-directed error mitigation, in particular, will be key to increasing gate depths and maintaining acceptable output fidelity. Pulse-level optimizations and specialized native gates will also be key enablers. Additionally, applications will be hybrid computations involving high-performance classical resources as well as quantum hardware serving as special-purpose accelerators. Effectively partitioning computations between these classical and quantum resources will be necessary to support realistic applications. Additionally, deep compiler optimization and classical simulation of Clifford and near-Clifford circuits can also be important classical investments towards more efficient quantum computations. Finally, defining abstractions that control compiler complexity yet selectively expose key physical machine properties will be also be a key area of research.

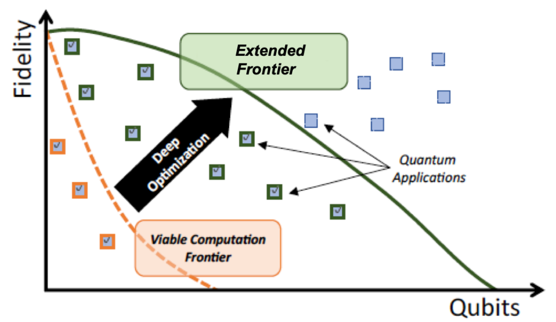

Fig. 1: Deep optimization advances the frontier of which quantum programs can be run successfully.

II. ERROR MITIGATION

Although gate error rates are improving, error mitigation techniques will be essential to allowing machines to support circuit depths with any potential of solving real problems. Since many error mitigation techniques depend upon qubit state and surrounding gate in the circuit, a compiler-directed approach can have significant benefits. Automated and semi-automated insertion of error mitigation techniques is desirable.

III. PULSE OPTIMIZATION

"Traditional" quantum compilation tools target a machine gate set, which is subsequently translated to control pulses for each individual gate to implement those gates. An interesting alternative is to translate application unitaries directly to pulses. Since the complexity of this translation scales exponentially with the size of the unitary, application programs must be broken into smaller blocks. Pulse techniques are also a powerful technique to implement specialized operations such as multi-qubit and qudit operators.

Current pulse optimization techniques suffer from poor scalability and poorly characterized hamiltonians. To address these, we need more efficient optimization methods, more robust solutions, and methods that use machines in the optimization loop. Borrowing from optimal control techniques from other fields such as robotics, trajectory-based and iterative-learning control techniques are promising directions along these lines.

IV. Hybrid Classical-Quantum Computation

With only a small number of quantum algorithms exhibiting promising advantages over classical computation, realistic applications of the foreseeable future will by necessity be hybrid combinations of classical computation and quantum kernels. In fact, applications may even use exponentially-scaling classical computation, as long as the compute time is practical and the advantage from the quantum kernel gives a total advantage for the entire computation. In the extreme, this is a contest of two exponentials, a (hopefully) exponential advantage for the quantum kernel and a potentially exponential cost to classical pre-processing, orchestration, and post-processing of the kernel code and data.

V. Deep Compiler Optimization

With small NISQ machines, compilers have often used relatively expensive classical computations to optimize quantum circuits, from mapping with look-ahead heuristics and SMT solvers, to gradient ascent for pulse optimization. As machines become larger, these optimizations can become considerably more expensive, yet some may still be worthwhile if the classical resources are available and the resulting quantum advantage is large enough. A precedent in the classical world is super-optimizing compilers, in which exhaustive search techniques are used to find optimal compilations for production binaries. While the quantum model is unlikely to be compile-once, run millions of times, the computational advantage of a quantum program may be days or years of classical compute time.

VI. Clifford-Guided Optimations

Many compiler optimizations depend upon program state and thus produce the best results when some dynamic program knowledge is available. In quantum systems, the lack of a ground truth for correct program results is an additional problem. We have found that approximating programs with all Clifford or mostly Clifford gates allows us to tractably simulate and learn about program behavior in ways that can significantly improve circuit fidelity.

Previous work has shown that Clifford approximations can be used to select which machines produce the best results for a particular circuit, and even reconstruct a higher fidelity results using multiple machines [12]. Preliminary results indicate that Clifford-based test circuits can predict compiler flag and optimization configurations that will give the best fidelity for a given application on a given machine. Beyond compilation, a Clifford-based approximation of a VQE ansatz can be used to give initial parameters that are orders of magnitude better than those given by Hartree-Fock [13].

There are also efforts to develop efficient classical simulation methods for near-Clifford circuits. These include a recent circuit-cutting approach [14] which, along with the IBM extended stabilizer simulator [15] and matrix-product-state simulation, can offer a suite of tools for near-Clifford simulation.

VII. Selectively Breaking Abstraction

As quantum machines scale, we will have to control classical compute time in our compilers and become more and more selective as to what low-level details to expose to our software and optimize for. A significant degree of modularity and abstraction will be necessary. The key is to identify the cross-layer optimizations that yield the highest fidelity benefits with the lowest classical compute complexity.

References

[1] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. Schuster, J. Majer, A. Blais, L. Frunzio, S. Girvin *et al.*, "Demonstration of two-qubit algorithms with a superconducting quantum processor," *Nature*, vol. 460, no. 7252, pp. 240–244, 2009.

[2] A. Zlokapa, S. Boixo, and D. Lidar, "Boundaries of quantum supremacy via random circuit sampling," *arXiv preprint arXiv:2005.02464*, 2020.

[3] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. Yavuz, T. Walker, and M. Saffman, "Observation of rydberg blockade between two atoms," *Nature Physics*, vol. 5, no. 2, pp. 110–114, 2009.

[4] T. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer *et al.*, "Multi-qubit entanglement and algorithms on a neutral-atom quantum computer," *Nature*, vol. 604, no. 7906, pp. 457–462, 2022.

[5] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland, "Demonstration of a fundamental quantum logic gate," *Physical review letters*, vol. 75, no. 25, p. 4714, 1995.

[6] A. Erhard, H. Poulsen Nautrup, M. Meth, L. Postler, R. Stricker, M. Stadler, V. Negnevitsky, M. Ringbauer, P. Schindler, H. J. Briegel *et al.*, "Entangling logical qubits with lattice surgery," *Nature*, vol. 589, no. 7841, pp. 220–224, 2021.

[7] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, "Even more efficient quantum computations of chemistry through tensor hypercontraction," *PRX Quantum*, vol. 2, no. 3, p. 030305, 2021.

[8] J. Sanders and E. Kandrot, *CUDA by example: an introduction to general-purpose GPU programming*. Addison-Wesley Professional, 2010.

[9] L. Dagum and R. Menon, "Openmp: an industry standard api for shared-memory programming," *IEEE computational science and engineering*, vol. 5, no. 1, pp. 46–55, 1998.

[10] A. B. Yoo, M. A. Jette, and M. Grondona, "Slurm: Simple linux utility for resource management," in *Job Scheduling Strategies for Parallel Processing: 9th International Workshop, JSSPP 2003, Seattle, WA, USA, June 24, 2003. Revised Paper 9*. Springer, 2003, pp. 44–60.

[11] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, "Focus beyond quadratic speedups for error-corrected quantum advantage," *PRX Quantum*, vol. 2, no. 1, p. 010103, 2021.

[12] G. S. Ravi, J. M. Baker, K. N. Smith, N. Earnest, A. Javadi-Abhari, and F. Chong, "Boosting quantum fidelity with an ordered diverse ensemble of clifford canary circuits," 2022.

[13] G. S. Ravi, P. Gokhale, Y. Ding, W. M. Kirby, K. N. Smith, J. M. Baker, P. J. Love, H. Hoffmann, K. R. Brown, and F. T. Chong, "CAFQA: A classical simulation bootstrap for variational quantum algorithms," in *ASPLOS*, T. M. Aamodt, N. D. E. Jerger, and M. M. Swift, Eds. ACM, 2023, pp. 15–29.

[14] K. N. Smith, M. A. Perlin, P. Gokhale, P. Frederick, D. Owusu-Antwi, R. Rines, V. Omole, and F. T. Chong, "Clifford-based circuit cutting for quantum simulation," in *50th ACM/IEEE Annual International Symposium on Computer Architecture, ISCA 2023, Orlando, Florida.*, 2023.

[15] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard, "Simulation of quantum circuits by low-rank stabilizer decompositions," *Quantum*, vol. 3, p. 181, sep 2019.

*FTC is an advisor to Quantum Circuits, Inc.

Not Only of IP Packets are Classical Networks Made Of: Adapting Knowledge from Other Domains to Scale Quantum Networks

Joaquin Chung[1]*, Neil Sinclair[2], Prem Kumar[3], Raj Kettimuthu[1]
[1]Argonne National Laboratory, [2]Caltech, [3]Northwestern University
email: chungmiranda@anl.gov, neils@caltech.edu, kumarp@northwestern.edu, kettimut@anl.gov

Topics: *quantum network models and architectures; codesign and integration*

1 Challenge

Humans make sense of abstract concepts by using metaphors. As we embark in the grand task of building large-scale quantum networks, the "classical" Internet has served as the major metaphor to reason about future quantum networks and eventually a quantum internet. For instance, an early proposal of a quantum internet stack directly maps quantum network functions into the TCP/IP stack [1]. However, using the classical Internet and specifically the TCP/IP stack as a blueprint for designing quantum networks can be misleading. In fact, Illiano et al. [2] have already warned about the impossibility of directly translating the TCP/IP stack into a quantum internet stack because of the marvels of quantum mechanics. Furthermore, the classical Internet in practice is more complex than the five layers of the TCP/IP stack. For instance, we have multiprotocol label switching (MPLS) serving as a "2.5" layer on carrier networks for more than 2 decades, similarly we have had encapsulation and tunneling running on the Internet for both security reasons and nefarious activities, and let's not forget about the cross-layering techniques used in ad hoc wireless networks to overcome noise in the channels. In this position paper we argue that rather than only using IP-based, packet switching networks as the blueprint for architecting future quantum networks, we should broader our view. In fact, the field of classical networking has produced a well of knowledge that can be adapted to the design and implementation of large-scale quantum networks. As Bruce Lee once said *"adapt what is useful, reject what is useless, and add what is specifically your own."* The rest of the paper presents the opportunities, our assessment, and the timeliness of our vision.

2 Opportunity

Quantum networks are reaching a level of maturity in which they have been demonstrated at campus and metropolitan scales [3, 4, 5]. Despite steps towards a quantum network stack, the diversity of physical solutions (i.e., qubit platforms, encoding schemes, and synchronization requirements) still demand for more specific (if not unique) quantum network implementations. Here we argue that the best lesson to learn from the classical Internet is not the TCP/IP stack but the hourglass shape of the stack itself. By hourglass shape we mean that IPv4 and IPv6 are the only protocols in the middle layer of the stack, while many more protocols exist in the upper and lower layers. If we consider entanglement generation and distribution the narrow waist of the hourglass (see Figure 1), we can accommodate the requirements of the heterogeneous physical platforms in the lower layers of the stack as well as the requirements of the applications above. We identify an opportunity for classical network engineers and researchers to work with physicists in the co-design of full-stack quantum networked systems that can operate autonomously as current classical networks

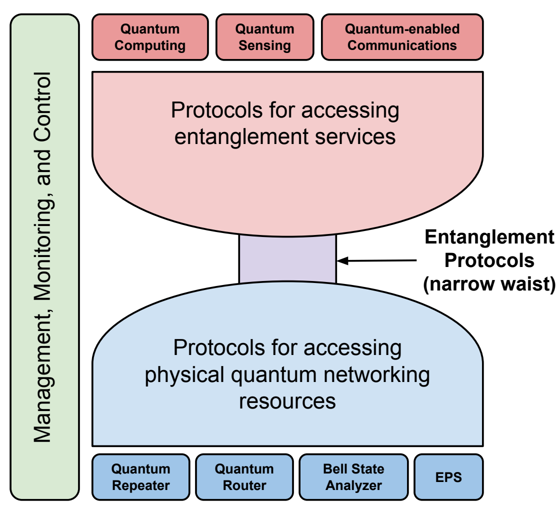

Figure 1: An hourglass-shaped quantum networking stack.

*Corresponding author: chungmiranda@anl.gov

do. This will promote an evolutionary pattern already seen in the early days of the classical Internet, where multiple protocols competed for the lower and upper layers of the stack. However, rather than competition, we encourage collaboration between groups for knowledge transfer and cross-pollination of ideas. Moreover, creating interfaces between these heterogeneous systems is a first step to scale quantum networks (towards a global quantum internet). Finally, we identify another opportunity to bring techniques from other classical networking domains (e.g., cellular networks, ad-hoc wireless networks, and the Internet of Things) that can help operationalize and scale quantum networks.

3 Assessment

A program that takes advantage of the opportunities presented in the previous section can be considered successful if the co-design of full-stack quantum networked systems produces demonstrations of quantum networks with the following features: (1) they can operate autonomously over extended periods of time, (2) they can run quantum computing or quantum sensing applications repetitively, and (3) they can server multiple users through friendly interfaces. Moreover, demonstrations of interoperability between such heterogeneous systems by using "quantum gateways" and inter-domain protocols can be considered the ultimate goal. To achieve these levels of success, we not only need to improve the performance of quantum devices, but we also need to develop classical or classical-quantum hybrid [6] control protocols and mechanisms for configuring, monitoring, and verifying quantum networks. This is of particular importance as protocols improve and evolve (e.g., single vs. two-photon schemes [7]), or introduction of advanced hardware implementations such as quantum frequency conversion, quantum-classical coexistence, or integrated electronics [8].

4 Timeliness

Several reasons make "now" the right time to start designing and evaluating full-stack quantum networked system: (1) quantum computing platforms are reaching commercial levels of maturity, (2) many quantum networking testbeds are being build across the world, and (3) quantum sensing is gaining more interest as a discipline. These events are making quantum computing, networking, and sensing more mainstream, which effectively attracts positive attention from engineers and practitioners that can help build the full-stack networked systems envisioned here. Moreover, the time is right to abstract interfaces and design classical control protocols that will help automate the operation of quantum networks and perform future interoperability tests. Success in this endeavor can be very impactful, propelling quantum networks from toys on a physics lab to useful tools that can benefit society as a whole.

References

[1] A. Dahlberg, M. Skrzypczyk, T. Coopmans, L. Wubben, F. Rozpundefineddek, M. Pompili, A. Stolk, P. Pawełczak, R. Knegjens, J. de Oliveira Filho, R. Hanson, and S. Wehner, "A link layer protocol for quantum networks," in *Proceedings of the ACM Special Interest Group on Data Communication*, SIGCOMM '19, (New York, NY, USA), p. 159–173, Association for Computing Machinery, 2019.

[2] J. Illiano, M. Caleffi, A. Manzalini, and A. S. Cacciapuoti, "Quantum internet protocol stack: A comprehensive survey," *Computer Networks*, vol. 213, p. 109092, 2022.

[3] M. Alshowkan, B. P. Williams, P. G. Evans, N. S. Rao, E. M. Simmerman, H.-H. Lu, N. B. Lingaraju, A. M. Weiner, C. E. Marvinney, Y.-Y. Pai, *et al.*, "A reconfigurable quantum local area network over deployed fiber," *arXiv preprint arXiv:2102.13596*, 2021.

[4] M. Pompili, C. D. Donne, I. te Raa, B. van der Vecht, M. Skrzypczyk, G. Ferreira, L. de Kluijver, A. J. Stolk, S. L. N. Hermans, P. Pawełczak, W. Kozlowski, R. Hanson, and S. Wehner, "Experimental demonstration of entanglement delivery using a quantum network stack," *npj Quantum Information*, vol. 8, oct 2022.

[5] J. Chung, E. M. Eastman, G. S. Kanter, K. Kapoor, N. Lauk, C. Peña, R. K. Plunkett, N. Sinclair, J. M. Thomas, R. Valivarthi, S. Xie, R. Kettimuthu, P. Kumar, P. Spentzouris, and M. Spiropulu, "Design and implementation of the illinois express quantum metropolitan area network," *IEEE Transactions on Quantum Engineering*, pp. 1–20, 2022.

[6] C. L. Cortes, P. Lefebvre, N. Lauk, M. J. Davis, N. Sinclair, S. K. Gray, and D. Oblak, "Sample-efficient adaptive calibration of quantum networks using bayesian optimization," *Physical Review Applied*, vol. 17, no. 3, p. 034067, 2022.

[7] M. Lucamarini, Z. L. Yuan, J. F. Dynes, and A. J. Shields, "Overcoming the rate–distance limit of quantum key distribution without quantum repeaters," *Nature*, vol. 557, no. 7705, pp. 400–403, 2018.

[8] D. Zhu, L. Shao, M. Yu, R. Cheng, B. Desiatov, C. Xin, Y. Hu, J. Holzgrafe, S. Ghosh, A. Shams-Ansari, *et al.*, "Integrated photonics on thin-film lithium niobate," *Advances in Optics and Photonics*, vol. 13, no. 2, pp. 242–352, 2021.

Quantum Readiness and Robustness Through Sustainable Quantum Software Practices

Daniel Claudino*, Elaine Wong, Sharmin Afrose, Peter Groszkowski, Justin Lietz, Francisco Rios, and Joseph Wang
Oak Ridge National Laboratory, Oak Ridge, TN, 37830, USA, *claudinodc@ornl.gov

Topic

Primary: Codesign and Integration; Secondary: Compilation

Challenge

The recent democratization of quantum computing is largely due to publicly available quantum hardware that can be used without extensive knowledge of the underlying physics and control electronics. This renders the gap between hardware and programmatic abstractions increasingly important and has led to a surge of mostly divergent quantum software endeavors,[1] and an uptick in software alternatives for quantum networks is starting to take place.[2] By far, these tools synergize with the platform-as-a-service (PaaS) model, now the standard method for accessing an increasing number of new quantum hardware architectures, which goes in the opposite direction of the current trend of edge computing and hinders more local approaches – e.g., Linux kernel driver model. These factors imply reduced lifetime, maintainability, and (re)usability, which are essential for sustainable software development,[3] and must be addressed to ensure the robustness of quantum software. The lack of proper documentation, design rationales, and specifications leads to entry barriers, reduce maintainability, and lower community engagement, eventually diverting resources to remedial and palliative tasks. By strengthening quantum compilation, codesign and integration techniques through the adoption of software sustainability practices, we can ensure that software tools enjoy a longer lifetime, and enable more effective collaboration across diverse systems and communities.

Opportunity

Software sustainability is an increasingly important topic[4,5] and in line with recent ASCR programs.[6] Its adoption can lead to a reduced demand for tangible resources (e.g., money, energy, and time) and more robust software tools. While quantum software is still in its infancy and is driven chiefly by prototypes and use cases, it presents an opportunity to incorporate sustainable practices into available tools and prioritize them from the onset of new projects. Doing so can help build a lasting future for quantum computing, thus sustainability should be embraced and permeate all quantum initiatives in which software plays a major role.

A simple definition of sustainability in software is its ability to function "agnostic of purpose".[7] One immediate way in which this can be accomplished is through refactoring of large codebases or monolithic applications into routines for specialized general tasks exposed via application programming interfaces (API). Furthermore, modularizing code following well established programming principles can further contribute to the sustainability of quantum computing software. Proper efforts in this direction increase the overall maintainability by making it easier to pin down errors in the software.

With quantum hardware continuing to function as accelerators, quantum software stacks are prone to demand much more flexible compilation strategies than classical software. In addition to cross-compilation to various hardware for portability, passes to encode the logical makeup of quantum error correction codes (QECC) will eventually be necessary. This complex landscape of "orthogonal" dimensions, such as native gate sets, hardware mapping and routing, QECCs, etc., highlights the need for tools in smaller, composable, plug-and-play-like modules in order to empower researchers to be effective at all levels of the stack.

Successfully expanding the scope, utility, efficiency, and robustness of software stacks for quantum computing and networking invariably depends on a proactive pursuit of interoperability, extensibility, and portability, ensuring that software is suitable and compatible with present and future hardware. The quantum intermediate representation (QIR),[8] based on the acclaimed LLVM compilation toolchain, embodies these principles and is a candidate to safeguard against ever-changing vendor software, while promoting a high degree of portability via the decoupling of software from specifics of the hardware, unifying quantum computing development efforts, and encouraging a more sustainable approach to software design.

Further progress can be achieved by promoting containerized computing (e.g., Docker), which frees end users from the need to customize applications for specific platforms. In addition to being more aligned with the demands of dependable software,[9] containers are also highly interoperable and a proven alternative to ensure reproducibility. Similarly, resource virtualization also promotes (re)usability by abstracting hardware details from higher levels in the stack, and together with containerization, can address the topic of portability. The benefits of these approaches will likely become more evident with the realization of embedding of quantum devices in leadership computing facilities and the growing need for efficient simulation tools for both quantum computing and networking to be portable into the emerging exascale regime.

Assessment

A current challenge in the literature is the adherence to an encompassing definition of sustainability and corresponding assessment metrics. Therefore, it is prudent to initially focus on simple and widely available parameters to gauge the effects of successful adoption of sustainable practices in quantum software. Such practices are expected to drive community engagement, thus proxy metrics of success may include: increased number of contributors and commits in version control systems (VCS), increased number of research artifacts that derive from sustainable quantum software (e.g., publications, presentations, and software dependencies), and indicators of frequent activity (statistics are easily obtained from VCSs). In the long run, it would also be possible to evaluate the effects on the lifetime of software projects. Sustainable software is apt to be more agnostic and easily adaptable to evolving hardware, resulting in its extended longevity and thus, usefulness.

Timeliness or Maturity

Quantum computing has long been viewed as a theoretical exercise with distant prospects. However, the recent explosion in better and larger quantum hardware has brought quantum software to the forefront, where it is now recognized as a critical bottleneck. Thus, careful software development is crucial to ensuring that we can reap all the potential benefits quantum computing has to offer. To tackle the challenges posed by this rapidly evolving landscape, durable and efficient software tools are necessary to enable research programs that depend on cross-platform benchmarking, circuit optimization, and post-processing data from a large number of qubits, to name a few. Continuously designing new tools from scratch with each hardware advancement is not only impractical but also unsustainable, and roadmaps from hardware vendors suggest that this software-hardware divide will only widen over time. Therefore, it is imperative to raise awareness of sustainable software development practices, including their integration into existing projects and their incorporation into future developments from the outset.

References

[1] Open-Source Quantum Software Projects. Accessed: 2023-04-18.

[2] Eight Quantum Networking Simulators. Accessed: 2023-04-26.

[3] C. Venters, L. Lau, M. Griffiths, V. Holmes, R. Ward, C. Jay, C. Dibsdale, and J. Xu. The Blind Men and the Elephant: Towards an Empirical Evaluation Framework for Software Sustainability. *J. Open Res. Softw.*, 2(1), 2014.

[4] C. Becker, R. Chitchyan, L. Duboc, S. Easterbrook, M. Mahaux, B. Penzenstadler, G. Rodriguez-Navas, C. Salinesi, N. Seyff, C. Venters, C. Calero, S. Akinli Kocak, and S. Betz. The Karlskrona manifesto for sustainability design, 2015. arXiv:1410.6968.

[5] C. Calero, M. Angeles Moraga, and F. Garcia. Software, Sustainability, and UN Sustainable Development Goals. *IT Professional*, 24(01):41–48, 2022.

[6] DOE ASCR Selects Seed Collaborations for Software Sustainability. Accessed: 2023-04-26.

[7] B. Penzenstadler, A. Raturi, D. Richardson, C. Calero, H. Femmer, and X. Franch. Systematic Mapping Study on Software Engineering for Sustainability (SE4S). In *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*, EASE '14, New York, NY, USA, 2014. Association for Computing Machinery.

[8] QIR Alliance. Accessed: 2023-04-18.

[9] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic Concepts and Taxonomy of Dependable and Secure Computing. *IEEE Trans. Dependable Secur. Comput.*, 1(1):11, 2004.

Software-Defined Quantum Networks

Connor Clayton Bobby Bhattacharjee Xiaodi Wu
University of Maryland, College Park
{cbclayto, bobby, xwu}@cs.umd.edu

Topic

Software-defined networking (SDN) in quantum networks: advantages, implementation, and scalability

Challenge

Quantum networks experience many of the same challenges faced by classical networks, such as routing and resource management. However, decoherence and the no-cloning theorem impose fundamental constraints that are unique to quantum networks: strict timing requirements and the inability (in general) to re-send quantum payloads mean that established store-and-forward or packet-switching techniques cannot be applied at the quantum layer. Instead, (at least early) quantum networks will follow a more circuit-switched model, wherein the "circuits" will be defined by sequences of entanglement-swapped "hops," leading to end-to-end transmission of qubits.

Fine-grained timing constraints and the closed-loop control structure required for circuit establishment and maintenance naturally points towards a somewhat centralized model of control, akin to how SDN controllers have been used for traffic engineering in classical networks. Several works [1, 2, 3] have already shown that centralized control admits substantial advantage over distributed control in entanglement-swapped networks under several different network models.

These protocols provide evidence for the benefits of SDN in quantum networks; however, the nature of the advantages SDN admits is not well understood. While SDN has been successful in the classical domain, it is unclear how this success will translate into the quantum domain and to what extent the advantages will be unique to quantum networks.

Additionally, methods of implementation for quantum SDN are not well understood, and current control schemes may limit scalability. State-of-the-art centralized protocols require all entanglement requests to route through a single controller and some [2, 3] require all nodes to send intermediate traffic to the controller on each round. This is in contrast with the implementation of OpenFlow [4]/NOX [5], where (ideally) only a small number of packets are routed through the controller since doing so is orders of magnitude slower than handling them on a switch.

Opportunity

A valuable tool and first step towards understanding and developing scalable centralized quantum routing protocols and network architectures will be a programming framework designed to support such designs. Classical SDN programming frameworks (e.g., OpenFlow [4]/NOX [5], Frenetic [6]) do not account for the specifics of quantum networks. For example, OpenFlow performs routing logic by pattern matching on packet headers, but there is no notion of packets or headers in entanglement swapping networks. *An SDN-framework designed bottom-up accounting for uniquely-quantum constraints could demonstrate the viability of SDN-control for quantum networks while also isolating deficiencies of centralized control.* If successful, such a framework could form the underlying platform for developing new quantum-specific protocols, and guide hardware design for deployable quantum networks.

The next step will be to consider the natural scalability limits of centralized SDN control mechanisms (in light of timing constraints imposed by quantum networks), and what techniques can be used to reduce control traffic in order to scale SDN-based deployment. Natural architectures to consider include hierarchies or networks of controllers. These constructions may be inspired by the design of classical software-defined networks, which are made practical by the observation that control need only be *logically* centralized, but may be *physically* decentralized. Such control structures may curtail controller overhead by distributing workload and locally centralizing computation.

Assessment

The utility of a quantum SDN programming framework lies primarily in its ability to illuminate properties of centralized control that are unique to quantum networks. This will involve identifying and demonstrating instances where centralized protocols improve (or fail to improve) network performance in the presence of constraints imposed by decoherence and the no-cloning theorem. Network performance may be measured using a variety of metrics, but these metrics will likely incorporate both throughput and fidelity (e.g., rate of entanglement generation meeting a fidelity threshold). The quality of an SDN protocol on a given network will be assessed via simulation; for example, using the quantum network simulator NetSquid [7].

The scalability of SDN architectures will also be assessed via simulation. Evaluations will analyze the behavior of existing and novel protocols at different network sizes through metrics including controller overhead (both traffic and time) and network throughput. A scalable protocol or architecture should have reduced overhead compared to existing work, yet comparable performance.

Timeliness or Maturity

Recent investment and work in the construction of quantum communication technologies has led to the emergence of several primitive quantum networks. Notable examples are located in Boston and Chicago domestically and the Netherlands and China abroad. *This nascent technology is approaching a stage where networks will become large, complex, and functional enough to benefit from more sophisticated control and management.*

SDN development has the potential to impact the quality of quantum networks by allowing strategic and therefore more effective use of quantum resources. Programming frameworks and a thorough understanding of the required controller architectures and associated protocols will allow quantum networks to scale more quickly by making fuller use of hardware as it develops.

References

[1] Shouqian Shi and Chen Qian. Concurrent entanglement routing for quantum networks: Model and designs. In *Proceedings of ACM SIGCOMM*, pages 62–75, 2020.

[2] Yangming Zhao and Chunming Qiao. Redundant entanglement provisioning and selection for throughput maximization in quantum networks. In *IEEE INFOCOM*, pages 1–10, 2021.

[3] Yangming Zhao, Gongming Zhao, and Chunming Qiao. E2e fidelity aware routing and purification for throughput maximization in quantum networks. In *IEEE INFOCOM*, pages 480–489, 2022.

[4] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. Openflow: enabling innovation in campus networks. *ACM SIGCOMM computer communication review*, 38(2):69–74, 2008.

[5] Natasha Gude, Teemu Koponen, Justin Pettit, Ben Pfaff, Martín Casado, Nick McKeown, and Scott Shenker. Nox: Towards an operating system for networks. *SIGCOMM Comput. Commun. Rev.*, 38(3):105–110, jul 2008.

[6] Nate Foster, Rob Harrison, Michael J. Freedman, Christopher Monsanto, Jennifer Rexford, Alec Story, and David Walker. Frenetic: A network programming language. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming*, pages 279–291, 2011.

[7] Tim Coopmans, Robert Knegjens, Axel Dahlberg, David Maier, Loek Nijsten, Julio de Oliveira Filho, Martijn Papendrecht, Julian Rabbie, Filip Rozpędek, Matthew Skrzypczyk, et al. Netsquid, a network simulator for quantum information using discrete events. *Communications Physics*, 4(1):164, 2021.

The Potential of Analog Quantum Computing

Carleton Coffrin and Marc Vuffray
Los Alamos National Laboratory
cjc@lanl.gov

Topic: Applications, Computing and programming models, Algorithms

Challenge: Although there has been some debate, Google's demonstration of *quantum advantage* on contrived computational tasks in 2019 [1] has been widely accepted as an important milestone in the progress of quantum computers. The next grand challenge in quantum computation is demonstrating that quantum computing devices can accelerate high-impact computations, a so-called *practical quantum advantage* [3]. It appears that such practical quantum advantage is starting to emerge in the realm analog quantum computation, which can have significant impacts on quantum computations related to scientific discovery.

It is widely agreed that the first applications where Quantum Computers (QC) will provide a transformational impact are in the simulation of quantum systems, such as modeling reaction mechanisms in chemistry [2]. Specifically, one can consider the time evolution of the following quantum dynamical system as suitable abstraction of this type of computational task,

$$i\frac{d}{dt}\left|\Psi(t)\right\rangle = H(t)\left|\Psi(t)\right\rangle ; \;\; t \in [0,T]; \;\; \left|\Psi_0\right\rangle \tag{1}$$

where the initial quantum state $|\Psi_0\rangle$ is evolved from time 0 to T under a potentially time varying Hamiltonian $H(t)$. For relatively small system sizes (i.e., ≈ 100 qubits), computing this type of dynamical evolution can be intractable to simulate with classical computers [3], presenting an appealing opportunity for quantum computers to support scientific questions that require these types of computations.

The dominant approach to developing quantum computer hardware has, thus far, been the so-called *gate-model*, which represents a quantum computation as a sequence of discrete gates applied to qubits. In this gate-based model, algorithms such as Totterization, Magnus Expansion and Quantum Signal Processing can be used to implement (1). However, the total steps (i.e., gate-depth) for such computations can easily exceed 10^6 for the smallest useful computations, leading to estimated computation times in the order of seconds to hours, under reasonable hardware assumptions [2]. This run time overhead presents a notable challenge for scaling this computational approach to simulation applications requiring tens of thousands of qubits or long evolution times (T). Consequently, without an algorithmic breakthrough, it appears that it will be many years before the gate-model of computation will be able to address important scientific computations taking the form of (1).

Opportunity: An important observation is that the gate-model is not the only model of quantum computation that is available to users of quantum computers. In many cases, quantum computing hardware has a lower-level interface where continuous time pules are used to drive the hardware's native quantum system [7]. This pulse-level model of the quantum computer follows the form of (1), where $|\Psi_0\rangle$ and $H(t)$ are constrained to some hardware-specific design. We refer to this type of quantum computation as *Analog Quantum Computing* (a.k.a., Quantum Simulators). The usefulness of Analog QC as a transitional technology in the NISQ era is well articulated by John Preskill in Section 6.10 of his seminal NISQ paper [8]. However, it now appears that this computational approach will have value beyond the NISQ era, due to the significant run time overheads required by fault tolerant quantum computation.

The core advantage of the Analog QC model over the gate-based model is a dramatic decrease is the computational resources needed to implement non-trivial quantum computations, both in qubit counts and run time. The primary challenge in Analog QC is identifying applications where the hardware's native Hamiltonian is a suitable match. The study of quantum magnetism provides one such good match for commercially available Analog QC hardware [5, 9, 4], but codesign research is required to better understand the possible matching of Analog QC architectures with impactful quantum simulation applications.

Assessment: The first measure of success of any QC simulation is to show that it can reproduce data collected from a physical experiment of a quantum systems (i.e., validation via a quantum testbed). The second measure of success is to show that the QC simulation can predict the outcome of a quantum experiment on models that have not been previously studied. Demonstrating both of these provides a strong indication that QC simulation is a suitable surrogate for the physical experiment. It is estimated that practically useful Analog QC hardware requires 100s to 1000s of qubits, which can be evolved for several microseconds of time. A programmable k-local Heisenberg model is likely sufficient for encoding Hamiltonians arising from a variety of applications, however more restrictive Hamiltonians can still be useful.

Timeliness and Maturity: In the last five years commercially available Analog Quantum Computing hardware (Pasqal, QuEra, D-Wave Systems, ...) has reached a point were it is starting to address important and realistic scientific computations. Some notable examples include: the study of quantum phase transitions in D-Wave's platform [6, 5] and emulation of realistic magnetic materials [4]; the identification of a topological spin liquid phase in QuEra's platform [9]; and promising optimization performance demonstrations using quantum computing hardware [10].

References

[1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[2] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoefler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vaschillo. Assessing requirements to scale to practical quantum advantage. *arXiv:2211.07629*, 2022.

[3] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller. Practical quantum advantage in quantum simulation. *Nature*, 607(7920):667–676, Jul 2022.

[4] A. D. King, C. D. Batista, J. Raymond, T. Lanting, I. Ozfidan, G. Poulin-Lamarre, H. Zhang, and M. H. Amin. Quantum annealing simulation of out-of-equilibrium magnetization in a spin-chain compound. *PRX Quantum*, 2:030317, Jul 2021.

[5] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud, E. Hoskinson, S. Huang, E. Ladizinsky, A. J. R. MacDonald, G. Marsden, R. Molavi, T. Oh, G. Poulin-Lamarre, M. Reis, C. Rich, Y. Sato, N. Tsai, M. Volkmann, J. D. Whittaker, J. Yao, A. W. Sandvik, and M. H. Amin. Quantum critical dynamics in a 5,000-qubit programmable spin glass. *Nature*, Apr 2023.

[6] A. D. King, S. Suzuki, J. Raymond, A. Zucca, T. Lanting, F. Altomare, A. J. Berkley, S. Ejtemaee, E. Hoskinson, S. Huang, E. Ladizinsky, A. J. R. MacDonald, G. Marsden, T. Oh, G. Poulin-Lamarre, M. Reis, C. Rich, Y. Sato, J. D. Whittaker, J. Yao, R. Harris, D. A. Lidar, H. Nishimori, and M. H. Amin. Coherent quantum annealing in a programmable 2,000 qubit ising chain. *Nature Physics*, 18(11):1324–1328, Nov 2022.

[7] D. C. McKay, T. Alexander, L. Bello, M. J. Biercuk, L. Bishop, J. Chen, J. M. Chow, A. D. Córcoles, D. Egger, S. Filipp, J. Gomez, M. Hush, A. Javadi-Abhari, D. Moreda, P. Nation, B. Paulovicks, E. Winston, C. J. Wood, J. Wootton, and J. M. Gambetta. Qiskit backend specifications for openqasm and openpulse experiments.

[8] J. Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, Aug. 2018.

[9] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):1242–1247, 2021.

[10] B. Tasseff, T. Albash, Z. Morrell, M. Vuffray, A. Y. Lokhov, S. Misra, and C. Coffrin. On the emerging potential of quantum annealing hardware for combinatorial optimization, 2022.

LA-UR-23-24475

Quantum-Centric Supercomputing

Thomas Alexander, Malcolm Carroll, Christopher Codella, Antonio Córcoles,* Ismael Faro, Jennifer R. Glick, Blake Johnson, Kevin Krsulich, Francisco Jose Martin Fernandez, Vincent R. Pascuzzi, Iskandar Sitdikov, and David Garcia Valinas
IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA
(Dated: May 1, 2023)

Topics: Models, Codesign and Integration

A. Challenge

Seamless communication and integration between layers of the cloud-based software stack, control hardware, and propagation of data are growing concerns as quantum systems scale up. Layered and optimized abstractions are needed to provide a frictionless and scalable user experience at multiple levels of expertise and for varied objectives. One layered abstraction, for example, is the error resilience layer. Error resilience strategies include the use of error mitigation [1], error suppression, and quantum error correction to overcome noisy hardware. A level above error resilience would include optimal circuit decomposition mapped to heterogeneous modular hardware. We emphasize a modular architecture, because modular hardware has become a leading paradigm to scale quantum processing units (QPU) and its supporting classical hardware [2]. Additional near and/or real-time classical processing may also be required for techniques like circuit knitting [3–5], that further enhances QPU capability. These layers are examples of the rapidly growing demands on circuit executions, classical compute, and the need to combine them seamlessly.

A primary challenge that arises is the orchestration of the classical resources [6] for these layers. The challenge includes forming optimal quantum circuit representations, management of the flows of data that span multiple timescales, and integration of operations with the QPU. This leads to related challenges such as doing the orchestration in a way that is accessible to a wide variety of users and implementing solutions in a scalable way.

Taking a historical perspective of state-of-the-art classical computing systems, high-performance computing (HPC) has evolved towards parallelized workflows as a natural way to increase speed. This happened by introducing a new technology, GPUs for example, for applications that could exploit complex, but simultaneous calculations. In a similar spirit, within the HPC framework, quantum systems can be envisioned as co-processors for certain applications by means of their own exclusive computing paradigm. This will require new programming models to allow the orchestration of classical and quantum workflows.

B. Opportunity

It now becomes natural to consider quantum processors as the next enhancement for HPC systems. A straightforward path to this integration relies on the effective leverage of HPC capabilities to support quantum workflows through orchestration, compilation, and different forms of output processing at various stages of the computation.

Incorporation of multi-cloud with HPC for workloads that require specialized classical compute may also enable parallel execution of quantum workloads. Serverless environments can relieve users from provisioning and maintaining compute infrastructure. Furthermore, at the next layer, dynamic orchestration of quantum and classical resources and workload composition, such as when using circuit knitting, can manage or eliminate the complexity of these two tasks on behalf of the user. Resource management could use policies that are either manually written or use machine learning models that monitor behavioral patterns of applications and components running in a quantum data center, becoming more effective over time through continuous training. For workload composition, orchestration includes problem analysis and optimization of circuit knitting operations either via developer-provided or selected tools and algorithms, or via automated analysis leading to optimize composition for user defined quality-of-service requirements.

Simplifying software frameworks, interfaces, and data structures to be more aligned with the hardware it will communicate with can reduce overheads in data transformations through the stack. Truly open source and communit-driven development of quantum-centric programming models [7], compilers [8], and computing frameworks would provide a common development interface and abstract away vendor specificity. That is, developing and maintaining bespoke software systems is an ongoing capital intensive process. Instead, collaborations among academia and industry can alleviate such hardships by offering a standardized computation model, while maintaining the ability of companies to leverage trade secrets through pluggable extensible mechanisms analogous to the successes of OpenCL, and other heterogeneous cross-platform abstractions. Adoption of open standards with significant community involvement is an ideal opportunity at this time [9].

Users can be more productive if they are not required to conform to an unfamiliar development environment that may include new programming languages and other tools to perform their work, or necessarily be concerned with

* Corresponding author: adcorcol@us.ibm.com

details of the underlying infrastructure. The programming model should be as simple as possible at the highest level but provide access to lower levels as needed, trading increased complexity for increased flexibility and customizability. At each level, default behavior and automation through declarative specification of program intentions may provide a completely specified system at the highest abstraction level, with open interfaces that allow a developer to substitute and tailor such behavior and automation at lower levels as necessary. Furthermore, adoption and establishment of open standards with significant community involvement is essential when defining interfaces and models at all levels.

C. Assessment

In general, we should see users obtaining improved runtimes, higher performance, stronger security, on increasingly larger workloads, which ultimately are beyond what classical computation alone could provide. Ultimately, success would mean that users such as model developers and industry experts are able to harness the value of quantum computation with software that hides underlying complexity without requiring expertise in quantum physics or in managing compute infrastructure. The programming model should be usable from within any development environment and enable integration with external commercial or proprietary software packages required to create enterprise-level workflows.

Success metrics for software architectures supporting quantum classical models depend on the application domain. A consistent set of standardized benchmarks can help assess the feasibility and then success of scaling-up existing system architectures (hardware and software) to yield practical and useful quantum computing. Along these lines, it will be relevant to identify how speed scales with increasing the number of qubits (e.g., increasing modules) analogous to HPC weak and strong scaling measures of the relative impact of the growth of other latencies with scale including how well cloud-accessed HPC resources are integrated. Ultimately, accelerated applications will motivate and justify coupling quantum computing with HPC resources and test its seamless integration with traditional scientific workflows.

D. Timeline and maturity

In recent years, there have been significant advances including the development of larger and more resilient quantum processors, improved techniques in error mitigation and circuit knitting, and quantum algorithms [10]. These advances have made it possible to explore industry-relevant use cases such as optimization, quantum chemistry simulations, and machine learning. Indeed, the continued maturation of workloads explored by industry places higher expectations on quantum hardware and will require increasingly specialized classical compute.

However, development of software for quantum computing must accelerate in preparation for scaled-up quantum hardware to match capabilities present today in classical enterprise and scientific computing. Now is the time to develop standardized abstractions and tooling to facilitate vendor development and user consumption of large-scale hardware capabilities. The level of sophistication and complexity of today's control systems and qubit technologies necessitates strong co-development.

This indicates the timeliness of designing and building full-stack architectures that can adequately address heterogeneous computation while also abstracting away complexity from the user. At the same time, the development of high-speed, low-latency communication networks, and the widespread adoption of cloud computing and virtualization technologies, make it possible to begin integrating classical and quantum computing resources seamlessly, and to orchestrate complex workflows.

[1] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Physical Review Letters **119**, 10.1103/physrevlett.119.180509 (2017).

[2] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, The future of quantum computing with superconducting qubits (2022).

[3] T. Peng, A. W. Harrow, M. Ozols, and X. Wu, Simulating large quantum circuits on a small quantum computer, Physical Review Letters **125**, 10.1103/physrevlett.125.150504 (2020).

[4] K. Mitarai and K. Fujii, Constructing a virtual two-qubit gate by sampling single-qubit operations, New Journal of Physics **23**, 023021 (2021).

[5] C. Piveteau and D. Sutter, Circuit knitting with classical communication (2022).

[6] R. Ranjan, B. Benatallah, S. Dustdar, and M. P. Papazoglou, Cloud resource orchestration programming: Overview, issues, and directions, IEEE Internet Computing **19**, 46 (2015).

[7] A. J. McCaskey, D. I. Lyakh, E. F. Dumitrescu, S. S. Powers, and T. S. Humble, XACC: A System-Level Software Infrastructure for Heterogeneous Quantum-Classical Computing (2019), arXiv:1911.02452 [quant-ph].

[8] T. M. Mintz, A. J. Mccaskey, E. F. Dumitrescu, S. V. Moore, S. Powers, and P. Lougovski, Qcor: A language extension specification for the heterogeneous quantum-classical model of computation (2019), arXiv:1909.02457 [cs.PL].

[9] A. Cross, A. Javadi-Abhari, T. Alexander, N. De Beaudrap, L. S. Bishop, S. Heidel, C. A. Ryan, P. Sivarajah, J. Smolin, J. M. Gambetta, and B. R. Johnson, Openqasmnbsp;3: A broader and deeper quantum assembly language, ACM Transactions on Quantum Computing **3**, 10.1145/3505636 (2022).

[10] A. R. Brodtkorb, C. Dyken, T. R. Hagen, J. M. Hjelmervik, and O. O. Storaasli, State-of-the-art in heterogeneous computing, Sci. Program. **18**, 1–33 (2010).

Practical Fault-Tolerant Quantum Computing Needs Intelligent Architectures

Poulami Das and Moin Qureshi

Georgia Institute of Technology, Atlanta, GA, USA

Email: {poulami,moin}@gatech.edu

1. INTRODUCTION

The high error-rates of quantum devices (about 1% on existing hardware [2]) limit us from running most practical quantum applications as they require lower error-rates (below 10^{-10}). Quantum error correction (QEC) bridges this gap by protecting quantum information. QEC codes encode a logical qubit by distributing information over redundant physical qubits. With increasing redundancy of the QEC code, the logical error-rate reduces exponentially if the physical error-rate is below a *threshold* [1]. Thus, by controlling the redundancy, QEC enables us to achieve the error-rate required to run any particular quantum application.

QEC codes use *data qubits* to store the quantum information and *parity qubits* to detect errors [10]. Fault-Tolerant Quantum Computers (FTQCs), as shown in Figure 1, ❶ use the control processor to send instructions to the qubits to execute *syndrome extraction circuits* that project errors on the data qubits onto the parity qubits, ❷ measure the parity qubits to obtain a bitstring of parity checks called *syndrome*, ❸ use decoders to analyze syndromes for identifying errors, and ❹ send the correction to the *control processor* so that errors are corrected in real-time. FTQCs enable computations by interleaving QEC cycles in between logical operations.

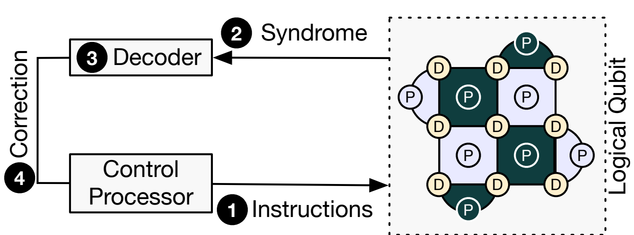

Figure 1: A logical qubit is encoded using *data (D)* and *parity (P)* qubits. Control processor sends instructions to extract a syndrome. Decoder uses syndromes to identify the correction and sends it back to the control processor.

2. CHALLENGES

Fault-tolerant quantum computing requires high bandwidth deterministic instruction supply to the quantum substrate as well as accurate, fast, and scalable error decoding. Unfortunately, only preliminary efforts have been made in the control processor design space, whereas most QEC studies have mainly resorted to offline decoding as software decoders are slow [6]. More recently, there has been a paradigm shift towards building real-time decoders using hardware. These designs typically optimize along three constraints: accuracy, latency, and scalability. For example, lookup table decoders are accurate and fast [5], but not scalable; approximate decoders are fast and scalable but sacrifice accuracy [4, 8]; fixed-function accelerator decoders are accurate, and fast, but not scalable [11]. Designing decoders that satisfy all three constraints is an open problem.

3. OPPORTUNITIES

With greater challenges come greater responsibilities. We propose a few concrete directions that could enable rapid and substantial in the field of fault-tolerant quantum computing.

Open-source tool-chains: We need software tool-chains to study the performance of QEC codes. Most tool-chains are the proprietary information of industrial groups with limited to no public access. Note that the time complexity of estimating the performance of QEC codes increases exponentially with increasing redundancy and lower device error rates. Although more recently, Google released their Stim [7] infrastructure, it's only a preliminary step, and more advanced software development frameworks are required.

Optimizing across dimensions: While surface codes are widely recognized as the most promising QEC code, they incur exponential overheads. Consequently, codes such as subsystem codes, Floquet codes, and QLDPC codes have been proposed that offer asymptotically lower overheads than surface codes at the expense of more irregular and denser device topology. Simultaneously, some qubit device topologies are more well-suited for denser connectivity compared to others. Designing accurate, high-performance, and scalable decoders that co-optimize over three dimensions: applications, qubit technologies, and QEC codes is important.

Real system studies and error models: Experimental demonstrations of QEC codes present new opportunities, insights, and challenges. For example, recent studies on Google Sycamore show how leakage errors could severely limit the performance of surface codes. While there exist device-level techniques to overcome these errors, their efficacy as well as overheads at the application level is unclear. There are similarly other sources of errors such as correlated errors. Developing decoding solutions that can tackle these errors will play a significant role in improving the performance of fault-tolerant systems. Architecture level solutions that adapt the code based on the required level of fault tolerant and that can dynamically reorganize the code based on the error events can greatly improve the efficacy of error correction while reducing the prohibitive overheads of fault tolerance.

Real-time control logic design: Control processors must ensure a high bandwidth deterministic instruction supply and adapt the QEC instruction schedules at runtime to revert the impact of errors. Failure to supply instructions and correct errors in real-time causes the QEC cycles to stall and consequently the qubits decohere. Unfortunately, most control processor designs are proprietary [3, 9], limiting research in this space. There are unique trade-offs in determining the location of the control processor- closer to the qubits at 4 Kelvin inside a dilution refrigerator means increased cooling overheads but reduced communication latencies, with lesser overheads at 77 Kelvin and room temperatures respectively. As cryogenic systems have very stringent power budgets, the control processors must be efficiently designed and potentially leverage the unique trade-offs of individual thermal domains to further optimize energy-efficiency.

Trade-offs in classical hardware development: Developing the control processor and decoders requires classical hardware development. While FPGAs present low-cost hardware development as well as flexibility and are being widely used in existing quantum platforms, their scalability and operational feasibility at low temperatures is an open problem. On the other hand, both cryogenic CMOS and superconducting designs are being considered for large-scale fault-tolerant systems (especially for superconducting systems). CMOS designs have large device densities and dense memory solutions while suffering from huge power dissipation. Superconducting designs on the other hand are extremely energy-efficient but lack the benefits of CMOS technology. Designing an efficient system-level architecture for control and QEC depends on the number of qubits in the system (dictates the complexity of the classical hardware), qubit device technology (determines the latency and bandwidth constraints), and the choice of the classical electronics.

4. ASSESSMENT

The metrics for success are somewhat well-known. The efficacy of quantum error-correction codes can be estimated using the metric of *Logical Error-Rate (LER)* and we would like to observe an exponential suppression of LER compare to the physical-error-rates as the code-distance is increased.

The short-coming in the modeling of error modalities (e.g. excluding leakage errors) would affect the LER of the code, as the LER observed on the real system would diverge from the theoretically expected LER. Thus, the metric of success of modeling would be in bridging the gap with the expected and observed LER on the real system.

The timeliness aspect of error decoding also impact on the effective LER observed on the system as taking a long latency for doing error decoding would cause timeout errors, causing increased rate of physical errors, beyond what can be corrected by the deployed QEC codes. Thus, we want the hardware to provide a decoding latency that achieves a LER close to the theoretical LER by minimizing timeout errors.

Finally, the figure-of-merit for the control processor and logic would be the number of qubits that can be supported by the given FPGA structure (to minimize the cost of custom ASIC logic) without leading to bandwidth bottlenecks or increased latency in supplying instructions (gates) to the qubits. Any delays would result in increased LER for the logical qubits as the likelihood of coherence errors is increased.

The ultimate goal would be to get to a logical qubit with LER initially of 10^-4 in the near term, then 10^-6 in the medium term, and finally 10^-8 in the long term. Achieving these goals will not only require a substrate with the right quality and quantity of qubits but also the right architectures to support these substrate with instructions. Finally, logical qubits needs to interact with other logical qubits, and will also need auxiliary structures such as T-factory to enable practical large-scale fault-tolerant quantum computing. This is a goal that will require collaboration with people across the stack, from devices, to information theory, to computer systems, and to algorithms and applications.

5. TIMELINESS AND MATURITY

The landscape of QEC is evolving rapidly with the introduction of newer QEC codes (such as QLPDC, floquet codes), a better understanding of newer fault modes (such as leakage, correlated errors), newer approaches to decoding (such as hardware-based, approximate techniques), and rapidly improving system sizes (now in the regime of up-to few hundreds of qubits). As QEC is essential to realize the true potential of quantum computers, there is increasing interest in studying QEC codes and real-time decoding. In recent years, several preliminary QEC experiments involving surface codes and subsystem codes have been successfully demonstrated.

REFERENCES

[1] D. Aharonov and M. Ben-Or, "Fault-tolerant quantum computation with constant error rate," *arXiv preprint quant-ph/9906129*, 1999.

[2] G. Q. AI, "Quantum computer datasheet," Accessed: June 19, 2021, https://quantumai.google/hardware/datasheet/weber.pdf.

[3] J. C. Bardin *et al.*, "29.1 a 28nm bulk-cmos 4-to-8ghz¡ 2mw cryogenic pulse modulator for scalable quantum computing," in *ISSCC*, 2019.

[4] P. Das *et al.*, "Afs: Accurate, fast, and scalable error-decoding for fault-tolerant quantum computers," in *HPCA*, 2022, pp. 259–273.

[5] P. Das *et al.*, "Lilliput: A lightweight low-latency lookup-table decoder for near-term quantum error correction," in *ASPLOS*, 2022.

[6] A. Fowler, "Towards sufficiently fast quantum error correction," 2017, conference QEC 2017. [Online]. Available: https://qec2017.gatech.edu

[7] C. Gidney, "Stim: a fast stabilizer circuit simulator," *Quantum*, vol. 5, p. 497, Jul. 2021. [Online]. Available: https://doi.org/10.22331/q-2021-07-06-497

[8] A. Holmes *et al.*, "Nisq+: Boosting quantum computing power by approximating quantum error correction," *arXiv preprint arXiv:2004.04794*, 2020.

[9] Intel, "Intel quantum," 2020. [Online]. Available: https://www.intel.com/content/www/us/en/newsroom/news/2nd-gen-horse-ridge-cryogenic-quantum-control-chip.html#gs.wc0iqx

[10] A. Y. Kitaev, "Fault-tolerant quantum computation by anyons," *Annals of Physics*, vol. 303, no. 1, pp. 2–30, 2003.

[11] S. Vittal *et al.*, "Astrea: Accurate quantum error-decoding via practical minimum-weight perfect matching," in *ISCA*, 2023.

Opportunities for quantum algorithmic research using nonlocal game models

Andrea Delgado [*], Kathleen E. Hamilton, Elaine Wong, Sarah Chehade and Jorge Ramirez Osorio
Oak Ridge National Laboratory, Oak Ridge, TN

Gabriel N. Perdue, Panagiotis Spentzouris
Fermi National Accelerator Laboratory, Batavia, IL

Topic: models, applications, algorithms, and codesign and integration

Quantum algorithms research has seen remarkable progress in recent years, fueled by advances in quantum computing (QC) hardware and a growing understanding of the unique properties and capabilities of quantum systems. However, we believe that quantum physics, especially entanglement, must be incorporated into the practice of algorithmic design in order to achieve a guaranteed quantum advantage. Entanglement is a fundamental quantum property, and we have yet to understand the full potential that it can bring into the context of information processing and algorithmic design. Nonetheless, entanglement is fundamentally hard to understand, and its analysis gives rise to challenging mathematical problems. To this end, exploring mathematical models can unlock the potential of QC and spur new approaches to developing quantum algorithms that enable the advancement of DOE science applications. **In this position paper, we advocate for nonlocal game (NLG) theory research. It is a promising opportunity to use mathematical models to understand and harness the potential of entanglement in developing quantum algorithms with a provable advantage for fundamental applications such as optimization and machine learning.**

TOPIC OVERVIEW — *The theory of NLGs bridges quantum algorithmic and mathematical research yielding robust tests for a quantum advantage that translates to practical use.* NLGs are mathematical models of a game with two or more players cooperating to achieve the best possible outcome without directly communicating with each other. Players receive questions from an external source (a referee) and must provide answers based on their knowledge. A quantum strategy is one in which the players share an entangled quantum state, thus correlating their actions. Otherwise, the strategy is classical. The 'value of the game' is the maximum success probability achievable by the players over a given strategy. Quantum advantage arises when the game value using a quantum strategy is larger than any classical strategy. Ideally, this would incorporate the fundamental properties of quantum information while adhering to the architectural decisions that bound system performance in practical applications. These applications include but are not limited to hidden linear functions (HLFs), graph state sampling (GSS), quantum networks, and hardware benchmarks.

There has been an increasing trend of interest in NLGs. In 2018, Bravyi et al.[3] showed that quantum advantage is achieved when solving the HLF problem by running parallel quantum algorithms in constant time due to quantum nonlocality. This result spawned the development of NLGs[6] to show a quantum advantage in solving the HLF problem using cluster states. Hasegawa and LeGall adapted NLGs to tackle the GSS problem, using expander graphs to handle corruption issues[10]. It was also shown around the same time that NLGs can provide an excellent benchmarking tool[7;14]. Bravyi et al.'s results were then strengthened to be applicable in different classes[1;2;8]. Furthermore, the quantum advantage is valid under circuit corruption in this setting[4;9]. Recently, NLGs were used to prove a remarkable result– the equivalence of two complexity classes which has astonishing consequences for operator algebra theory by resolving a long-standing conjecture[12].

CHALLENGES — We identify three challenges in designing quantum algorithms based on NLGs: 1) identifying algorithms with a provable quantum advantage, 2) establishing proof that this advantage is not susceptible to dequantization, and 3) that these algorithms are useful in practice.

[*] delgadoa@ornl.gov

The first challenge is understanding the limits of quantum algorithms and the conditions under which they offer advantages over classical approaches. This remains an active area of research, including the algorithms' computational complexity and investigating the consequences of noise and decoherence in practical implementations. In addition, understanding how to incorporate and maximize the potential of entanglement in information processing and algorithmic design is still an open question. Symmetry is another property that has been incorporated into quantum algorithm design, but its full utility has yet to be completely established. Expanding the NLG formalism to more general problem structures[13] may further illuminate symmetry's role in quantum search algorithms[11]. The second challenge is that the quantum advantage frontier is pushed further by efforts at dequantization, first exhibited in the case of recommendation systems[15]– and then extended to solutions of low-rank linear systems[5], to devise quantum-inspired algorithms with similar performance to their quantum counterparts. The third challenge is finding problems that have a quantum advantage and are robust against noise and dequantization and relating them to useful problems. For example, a perfect game strategy has been established to be undecidable for specific game strategies and algorithms without practical use[12].

OPPORTUNITY — NLGs provide an excellent example of tasks that benefit from manipulating quantum resources and have applications across many fields of study that intersect with quantum information science, including cryptography, computer science, and other mathematical areas. In addition, NLGs are a powerful way (in terms of separation of classes) to demonstrate a quantum advantage that, in some instances, is robust to dequantization. Finally, the states underlying NLGs, and particularly highly-entangled states, have known connections to error-correcting codes, which may connect to interesting algorithmic algebraic problems found in scientific computing.

ASSESSMENT — The long-term success metric is the development of solid mathematical proof or results that indicate a family of NLGs with a noise-robust quantum advantage that is impervious to dequantization and can have practical applications in areas of relevance to the DOE. Near-term success focuses on benchmarking or connecting predicted performance under corrupting noise[4;10] to hardware performance. A longer-term assessment would include incorporating the measures of GSS[10] or HLF[3] into real-world problems. For perspective, the GSS problem generalizes the HLF problem. Instead of searching for a specific bitstring in a graph represented as a 2D grid, the GSS problem considers finding arbitrary bitstrings on expander graphs. The HLF and GSS problems can have applications for quantum searches executed over graphs.

TIMELINES OR MATURITY — NLGs have been played on near-term hardware, and we expect current limitations to be resolved as hardware capabilities advance and mid-circuit and conditional measurements become available. Bridging the mathematical theory of NLGs and the design of practical quantum algorithms presents an excellent opportunity for co-design and collaboration between existing DOE quantum testbeds and researchers in quantum computer science and algorithms.

References: [1] Bene Watts, A., et al. In Proc. 51 Ann. ACM Symp. Theory Comput. (New York, NY, USA, 2019), STOC 2019, Assoc. Comp. Mach., p. 515–526. [2] Bene Watts, A., et al., 2023. [3] Bravyi, S., et al. Science 362, 6412 (10 2018), 308–311. [4] Bravyi, S., et al. Nat. Phys. 16 (7 2020), 1040–1045. [5] Chia, N. H., et al., 2018. [6] Daniel, A. K., et al. Phys. Rev. Res. 4 (7 2022), 033068. [7] Daniel, A. K., et al., 2022. [8] Grier, D., et al. In Proc. 52 Annu. ACM Symp. Theory Comput. (New York, NY, USA, 2020), STOC 2020, ACM, p. 875–888. [9] Grier, D., et al., 2021. [10] Hasegawa, A., et al. In 32nd Int'l Symp. on Alg. and Comp. (ISAAC 2021) (Dagstuhl, Germany, 2021), H.-K. Ahn and K. Sadakane, Eds., vol. 212 of Leibniz Int'l Proc. in Inform. (LIPIcs), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, pp. 74:1–74:16. [11] Janmark, J., et al. Phys. Rev. Lett. 112 (5 2014), 210502. [12] Ji, Z., et al. Commun. ACM 64, 11 (2021), 131–138. [13] Kravchenko, D., et al. In Reachability Problems: 14th International Conference, RP 2020, Paris, France, October 19–21, 2020, Proceedings (2020), Springer, pp. 83–98. [14] Sheffer, M., et al. Adv. Quantum Technol. 5, 3 (2022), 2100081. [15] Tang, E. In Proc. 51 Ann. ACM Symp. Theory Comput. (6 2019), ACM.

Tailoring Quantum Error Correction Codes to Sparse Hardware: Compiler and Software Innovations

Topic: Quantum Compiler, Quantum Error Correction.
Yufei Ding, Assistant Professor, University of California, Santa Barbara.

Challenge: Quantum error correction (QEC) serves as the fundamental component of fault-tolerant quantum computation; however, there often exists a structural mismatch between QEC codes and the underlying hardware. For instance, the surface code[1] necessitates a 2D-lattice qubit structure, with each qubit coupled to four neighbors. Unfortunately, many cutting-edge quantum processors, such as IBM's heavy-hexagon-architecture chip[2] and Rigetti's octagonal-architecture device[3], do not readily support such an architecture.

On one hand, compiling a general quantum circuit onto a sparse quantum device has been well studied [4, 5] in recent years. However, general quantum compilers are not suitable for compiling QEC codes. Firstly, they don't distinguish data qubits from other qubits. They may move data qubits frequently, invalidating logical operations designed for a fixed data qubit layout [1]. Secondly, the SWAP gates they use to overcome the connectivity issue make the compiled measurement circuits too error-prone for practical error correction. Finally, they do not account for the constraints on the syndrome measurement circuits [1], e.g., the order of CNOT gates between syndrome qubits and data qubits.

On the other hand, certain manual efforts have been undertaken to investigate the integration of quantum error correction (QEC) codes and hardware accelerators, such as the manual assignment of data qubits to these accelerators[6]. While this approach is manageable on a small scale, it becomes increasingly untenable to manually devise code mapping protocols as code sizes and hardware platform complexities expand. Furthermore, due to the common presence of hardware imperfections and variations in the fabrication and execution processes, designing and deploying quantum error correction codes necessitates increasingly specialized skills and expertise.

Opportunity: Our research group has developed a QEC compiler framework [7] that can generalize surface codes to any sparse superconducting hardware, without the need for a 2-D lattice hardware topology. To the best of our knowledge, this work is the first attempt that formalizes the surface code synthesis problem. In particular, we identify three key challenges of the surface code synthesis: data qubit allocation, bridge tree construction, and stabilizer measurement schedule. The proposed framework is the first automated solution to the surface code synthesis problem, as far as we know. We also show that the surface code synthesized by our framework can achieve comparable or even better error correction capability than manually designed QEC codes. This preliminary study is specifically designed for the surface code.

This preliminary study is specifically designed for the surface code. To address the challenge of general QEC code to hardware mapping, we suggest two research opportunities that could potentially build upon each other. First, we may leverage MaxSAT to model the constraints in the mapping process and generate an optimal mapping for the given QEC code and hardware platform. This approach allows us to capture the complexity of the mapping problem and provide an optimal solution while taking into account the unique constraints of each hardware platform. However, this approach may become computationally intractable for larger QEC codes and hardware systems. To overcome this limitation, we further suggest to develop a machine learning-based approach that leverages the insights and knowledge gained from the MaxSAT approach to create a more scalable and modular solution. By combining the two approaches, we can provide a robust and efficient solution for QEC code to hardware mapping for current and future quantum computing systems.

Assessment: A crucial metric for evaluating potential solutions for compiling QEC codes is the required quantum resources, particularly the number of physical qubits needed to compose two logical qubits while achieving a logical error rate between 10^{-6} and 10^{-9} with

QEC protection. Detailed success metrics may include improvements in error rates, qubit usage, and overall quantum resource requirements. Estimates or projections of required quantum resources should consider the complexity of the target hardware, QEC code size and type, and expected error rates while remaining flexible to accommodate emerging and evolving hardware architectures. The performance of both the MaxSAT-based compiler and the ML-based compiler will also be evaluated on various benchmark QEC codes and hardware platforms to demonstrate their effectiveness and efficiency.

Timeliness or maturity: This research topic is timely and steadily maturing. On the hardware side, quantum technology has made remarkable progress over the past decade, culminating in the first demonstration of quantum supremacy in 2020 [8]. Recent hardware advancements include significant improvements in physical error rates, single-qubit addressability, and manufacturing scalability. On the QEC code side, groundbreaking work by Google [9] and Yale's team [10] has achieved two critical milestones in quantum computing: enhanced QEC power with more physical qubits, and the development of a QEC-protected logic qubit that outperforms even the best isolated physical qubit.

However, the race is far from over. We envision the next milestone in the race is to construct a quantum system with at least two logical qubits and to reach a logic error rate ranging from 10^{-6} to 10^{-9} with QEC protection. This milestone is essential because practical applications like integer factoring would require 10^6 to 10^9 gates to outperform classical computing. Success in this area would lead to increased confidence in the long-term viability and practicality of quantum computing.

REFERENCES

[1] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland. 2012. Surface codes: Towards practical large-scale quantum computation. *Physical Review A* 86, 3 (2012), 032324.

[2] Eric J Zhang, Srikanth Srinivasan, Neereja Sundaresan, Daniela F Bogorin, Yves Martin, Jared B Hertzberg, John Timmerwilke, Emily J Pritchett, Jeng-Bang Yau, Cindy Wang, et al. 2020. High-fidelity superconducting quantum processors via laser-annealing of transmon qubits. *arXiv preprint arXiv:2012.08475* (2020).

[3] Alysson Gold, Anna Stockklauser, Matt Reagor, Jean-Philip Paquette, Andrew Bestwick, Cody James Winkleblack, Ben Scharmann, Feyza Oruc, and Brandon Langley. 2021. Experimental demonstration of entangling gates across chips in a multi-core QPU. *Bulletin of the American Physical Society* (2021).

[4] Gushu Li, Yufei Ding, and Yuan Xie. 2019. Tackling the qubit mapping problem for NISQ-era quantum devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 1001–1014.

[5] Swamit S Tannu and Moinuddin K Qureshi. 2019. Mitigating measurement errors in quantum computers by exploiting state-dependent bias. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. 279–290.

[6] Laird Egan, Dripto M Debroy, Crystal Noel, Andrew Risinger, Daiwei Zhu, Debopriyo Biswas, Michael Newman, Muyuan Li, Kenneth R Brown, Marko Cetina, et al. 2020. Fault-tolerant operation of a quantum error-correction code. *arXiv preprint arXiv:2009.11482* (2020).

[7] Anbang Wu, Gushu Li, Hezi Zhang, Gian Giacomo Guerreschi, Yufei Ding, and Yuan Xie. 2022. A synthesis framework for stitching surface code with superconducting quantum devices. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*. 337–350.

[8] F. Arute, K. Arya, R. Babbush, D. Bacon, J. Bardin, R. Barends, R. Biswas, S. Boixo, F. Brandão, D. Buell, B. Burkett, Y. Chen, Zijun Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, Keith Guerin, Steve Habegger, M. Harrigan, M. Hartmann, A. Ho, M. Hoffmann, Trent Huang, T. Humble, S. Isakov, E. Jeffrey, Zhang Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, Mike Lindmark, E. Lucero, Dmitry I. Lyakh, Salvatore Mandrà, J. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Niu, E. Ostby, A. Petukhov, John C. Platt, C. Quintana, E. Rieffel, P. Roushan, N. Rubin, D. Sank, K. Satzinger, V. Smelyanskiy, Kevin J. Sung, M. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. Yao, P. Yeh, Adam Zalcman, H. Neven, and J. Martinis. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574 (2019), 505–510.

[9] 2023. Suppressing quantum errors by scaling a surface code logical qubit. *Nature* 614, 7949 (2023), 676–681.

[10] VV Sivak, A Eickbusch, B Royer, S Singh, I Tsioutsios, S Ganjam, A Miano, BL Brock, AZ Ding, L Frunzio, et al. 2022. Real-time quantum error correction beyond break-even. *arXiv preprint arXiv:2211.09116* (2022).

Quantum computing for modeling of classical systems

I. Y. Dodin,[1,2,*] I. Novikau,[1,†] and E. A. Startsev[1,‡]
[1]*Princeton Plasma Physics Laboratory, Princeton, NJ 08543, USA*
[2]*Department of Astrophysical Sciences, Princeton University, Princeton, New Jersey 08544, USA*

Quantum computing is considered to be a promising way to speedup simulations of quantum systems. Can it also help solve classical-physics problems? Recent studies indicate that this may be feasible, but many challenges remain, and addressing them requires more than toy-problem solving with ad hoc tricks. Instead, theory's focus should be shifted to generic methods and practicality.

Topic: applications, algorithms

Challenge: One of the important potential applications of quantum computing (QC) is modeling of physical systems. Typically considered as candidates for such modeling are quantum N-body systems with $N \gg 1$, which have notoriously large configuration spaces and are impossible to simulate on classical computers from first principles. However, modeling of classical physics could also benefit from QC for the same reasons. This is especially the case for multi-scale dynamics (e.g., fluid turbulence) or dynamics of macroscopically many free particles whose individual motions need to be resolved (e.g., in plasma problems). High-fidelity simulations of such dynamics are critical for various practical applications (e.g., nuclear fusion) but, even with modern supercomputers, remain limited to specific regimes or involve severe approximations for the lack of better options. Thus, figuring out how to model classical systems on quantum computers is an important problem to solve.

However, practical classical problems can be extremely challenging to map on (gate-based) quantum computers, which are naturally suited to perform only linear unitary operations. The specific reasons are as follows:

- Classical problems of interest are usually nonlinear.
- The rare classical problems that are of interest *and* linear are typically non-conservative and therefore cannot be solved using straightforward QC techniques such as Hamiltonian simulations.
- The rare classical problems that are of interest, linear, *and* conservative often have non-sparse Hamiltonians. They are difficult to implement in quantum circuits and lack simple properties that would allow for particularly efficient quantum algorithms.

Opportunity and assessment: Figuring out how to deal with the issues listed above is a research opportunity that is yet to be seized. We explored it, preliminarily, in [1–3] (see also the recent review [4]), and our assessment of the current challenges is as follows.

* idodin@pppl.gov
† inovikau@pppl.gov
‡ estarts@pppl.gov

Hamiltonian simulations. — Let us start with Hamiltonian simulations, which can be used, for example, for modeling linear waves in collisionless plasmas [1, 2, 5]. In this case, one needs to find a system representation in which system's state vector ψ (e.g., rescaled electromagnetic field) is governed by the Schrödinger equation $\mathrm{i}\partial_t\psi = \hat{H}\psi$ with Hermitian $\hat{H}$. Then, one needs to encode the initial state ψ_0 and the Hamiltonian $\hat{H}$ as an oracle, develop a circuit approximating the unitary propagator $\hat{U}_t = \exp(-\mathrm{i}\hat{H}t)$ (assuming that $\hat{H}$ is independent of time t) to calculate $\psi(t) = \hat{U}_t\psi_0$, and eventually perform measurements. None of these steps are trivial.

In particular, encoding of a classical Hamiltonian (which is typically non-unitary) with a unitary circuit requires sophisticated techniques such as the block encoding [6]. Notably, efficient construction of the corresponding oracles cannot be guaranteed, especially considering that a classical $\hat{H}$ can be non-sparse [1, 5] and usually is not a direct product of Pauli matrices, unlike typical Hamiltonians in quantum mechanics. In simple problems, the corresponding oracles can be constructed manually ad hoc and exponential quantum speedup can be achieved [2]. However, machine-learning-based algorithms will likely be necessary to perform the encoding of Hamiltonians for large-scale simulations in the future.

Encoding of $\hat{U}_t$ for a given $\hat{H}$ is not trivial either. Simple trotterization [7] is non-optimal, because a classical $\hat{H}$ is not efficiently represented through Pauli matrices. (A brute-force decomposition of such $\hat{H}$ into Pauli gates generally makes the circuit depth scale as a polynomial of the matrix size.) Instead, the so-called quantum signal processing (QSP) method was proposed, which involves qubitization, i.e., constructing $\hat{U}_t$ out of $\hat{H}$ via a sequence of rotations [8, 9]. However, implementation of the complex QSP circuits might require availability of elementary gates controlled by multiple qubits. (Having other customized gates may also help [10].) Otherwise, emulation of such gates with simpler gates will be necessary, which will increase the circuit size polynomially with the number of controlled nodes [11].

Non-Hermitian Hamiltonians, matrix inversion. — To deal with non-Hermitian Hamiltonians in linear problems, one can replace the initial-value problem with a boundary-value problem of the form $\hat{A}\psi = b$, where the matrix $\hat{A}$ and the vector b are given and ψ is the vector that needs to be found. Problems like these emerge, for example, in modeling plasma heating with linear radiofrequency waves launched by external antennas [3]. In such applications, the matrix $\hat{A}$ is typically non-unitary, so, again, block encoding is necessary to represent it with a unitary circuit, and similar concerns as for Hamiltonian simulations apply. The matrix proportional to $\hat{A}^{-1}$ can be calculated using various methods. As of now, particularly promising is the state-of-the-art method called quantum singular value transformation (QSVT) [12, 13], which is a generalization of the QSP for non-Hermitian matrices. However, interesting classical problems are usually ill-conditioned, which significantly reduces the QSVT efficiency [3]. Also, the rotation angles necessary for the construction of QSVT circuits must be precalculated classically, and how to do this for large-scale ill-conditioned problems remains to be understood. Such application will likely require preconditioners [14] that will reduce the condition number of $\hat{A}$ while remaining simple enough to be block-encoded.

Initial-value problems with non-Hermitian Hamiltonians may be solvable too, specifically, using non-unitary gates [15] or the imaginary-time algorithm [16]. However, the latter requires tomography of quantum states, which significantly reduces the quantum speedup, and overall, these possibilities remain to be explored.

Nonlinear problems. — Nonlinear problems, which are most interesting for classical applications, cannot be solved on a quantum computer directly. Various ways have been proposed to deal with this nuisance, and most of them involve embedding of nonlinear systems into linear ones, either through quantization or other methods; see, e.g., [1, 17–20]. (Another possible approach is hybrid quantum–classical algorithms [21].) This significantly increases the system dimension and may completely preclude a quantum speedup. Hence the question: how can one truncate linear embeddings efficiently such that the system dimension would remain manageable?

Some embeddings can be efficient for toy models yet inapplicable for simulating of anything interesting [22]. For example, the well-known Carleman embedding is useful only when nonlinearity is weaker than dissipation; e.g., in fluid dynamics, the Reynolds number must be sufficiently small [20]. One alternative that seems promising is the so-called Koopman–von Neumann approach [1, 17], where the nonlinear evolution of a system is described in terms of the linear evolution of the corresponding probability distribution in phase space. Compared to the Carleman embedding, this approach is more robust, albeit also more demanding computationally. In any case, to our knowledge, no nonlinear classical dynamics beyond toy models has been simulated with any quantum algorithm so far. Thus, the jury is still out on how such simulations will be carried out in the future.

Timeliness: Recent studies indicate that modeling of classical systems may be an area where QC can make a difference, in terms of both physics and practical applications. However, this is also an area where toy problems are irrelevant; they can be solved already on existing computers and therefore do not require QC. For QC to become a useful and functional tool for classical physics (plasma physics, hydrodynamics), focus should be made on systematic theory, generic methods, and practicality rather than ad hoc tricks and theorem proving for algorithms that will never be applicable to nontrivial problems. So far, little attention has been paid to this subject by the QC community, as also seen from the laundry list of the outstanding issues discussed above. One can consider this as a sign that quantum simulations of classical systems will never be practical. Or the conclusion can be that the field is full of ripe opportunities for those who are up to the challenge. Which way is it going to be?

This paper is based on the research conducted under the Laboratory Directed Research and Development (LDRD) Program at Princeton Plasma Physics Laboratory, a national laboratory operated by Princeton University for the U.S. Department of Energy under Prime Contract No. DE-AC02-09CH11466.

[1] I. Y. Dodin and E. A. Startsev, Phys. Plasmas **28**, 092101 (2021), arXiv:2005.14369 (2020).
[2] I. Novikau, E. A. Startsev, and I. Y. Dodin, Phys. Rev. A **105**, 062444 (2022).
[3] I. Novikau, I. Y. Dodin, and E. A. Startsev, arXiv:2212.09113, to appear in Phys. Rev. Appl..
[4] I. Joseph *et al.*, Phys. Plasmas **30**, 010501 (2023).
[5] A. Engel, G. Smith, and S. E. Parker, Phys. Rev. A **100**, 062315 (2019).
[6] A. M. Childs and N. Wiebe, Quantum Info. Comput. **12**, 901 (2012).
[7] M. Suzuki, Phys. Lett. A **146**, 319 (1990).
[8] G. H. Low and I. L. Chuang, Phys. Rev. Lett. **118**, 010501 (2017).
[9] G. H. Low and I. L. Chuang, Quantum **3**, 163 (2019).
[10] Y. Shi *et al.*, Phys. Rev. A **103**, 062608 (2021).
[11] A. Barenco *et al.*, Phys. Rev. A **52**, 3457 (1995).
[12] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, in Proc. 51st Annual ACM SIGACT Symp. Theory of Computing (Assoc. Computing Machinery, NY, 2019), p. 193.
[13] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, PRX Quant. **2**, 040203 (2021).
[14] J. Golden, D. O'Malley, and H. Viswanathan, Sci. Rep. **12**, 22285 (2022).
[15] E. Zapusek, A. Javadi, and F. Reiter, Quantum Sci. Technol. **8**, 015001 (2022).
[16] M. Motta *et al.*, Nature Phys. **16**, 205 (2020).
[17] I. Joseph, Phys. Rev. Res. **2**, 043102 (2020).
[18] A. Engel, G. Smith, and S. E. Parker, Phys. Plasmas **28**, 062305 (2021).
[19] S. Lloyd *et al.*, arXiv:2011.06571.
[20] J.-P. Liu *et al.*, Proc. Nat. Acad. Sci. **118**, e2026805118 (2021).
[21] M. Lubasch *et al.*, Phys. Rev. A **101**, 010301(R) (2020).
[22] R. F. S. Andrade and A. Rauh, Phys. Lett. A **82**, 276 (1981).

Hybrid Matter-Photon Modules for Distributed Quantum Computing and Networks

Shengwang Du

Department of Physics, University of Texas at Dallas, Email: dusw@utdallas.edu

Topic: Distributed quantum computing and quantum networks

Challenge

Quantum computing and networking have not yet met each other. Although now quantum computers have entered industry, their remote connections are still classical due to the lack of efficient conversion interface between local computing qubits and flying photonic qubits. On the other side, experimental studies of quantum networks mainly focus on entanglement distribution and security communication, with little in networking remote quantum computers.

Why do we need distributed quantum computing and networking multiple quantum processors? It is because the size of a local quantum computer is always limited by physical constrains. An ideal quantum computer of N>300 qubits is powerful because its 2^N dimensional Hilbert space cannot be sorted even by using up all atoms in the universe as classical bits. However, in a real quantum computer, there is a limitation of how many qubits can be directly interacted. *The computational parallelism and power of a quantum computer is determined by the **dimension** of its qubits Hilbert space which can be simultaneously addressed* (through single qubit operations and measurements) *and parallelly controlled* (through control gate operations). For example, for a quantum computer with N qubits, if a control gate operation cannot be performed between any two arbitrary qubits, its computing power is $< 2^N$. Therefore, number of qubits in a quantum computer is not a "precise" measure of power, though it has been commonly used to catch public eyeballs. For a superconductor quantum computer, its real computing power is always much lower than the claimed because a superconductor qubit can only interact with its nearby qubits and connecting two distant qubits must go through many sequential gate operations. For this reason, although IBM has built a 433-qubit processor called Osprey in 2022 and plans to build a 1121-qubit device by 2023 [1], they are all still far away from beneficial commercialization. ColdQuanta, the leading neutral atom quantum computer company, is targeting a 1,000-qubit quantum computer called Hilbert in 2024, but its qubit interaction connectivity is only at 4:1 [2]. Trapped-ion-based quantum computers [3] have a similar challenge.

Connecting two N-qubit remote quantum computers classically, the dimension of their combined Hilbert space is only $2\times 2^N=2^{(N+1)}$. If two N-qubit quantum computers are fully connected though quantum links, the dimension of their combined Hilbert space could reach $2^{(2N)}$ which is much more powerful than two independent quantum computers. To network two remote quantum computers with photonic links to extend their joint Hilbert space beyond the classical connection requires efficient conversion between local computing qubits and flying photonic qubits. However, there has been no such modules for existing quantum computing platforms. This is due to the fact that a qubit based on a single atom or ion (or artificial atom such as superconducting circuit or quantum dot) has a very small cross section ($\sim\lambda^2$) to interact with a single photon. One may argue to increase the coupling by loading atomic qubits into a high fineness cavity, which is possible for one qubit, but it is impossible to load more than 100 single-atom qubits without losing their independency because a single-mode cavity photon would interact with all qubits simultaneously.

Another problem of the existing quantum computing platforms is their decoherence time which limits circuit depth or computation steps. Decoherence reduces a pure quantum state to a mixed state and the quantum computing must stop to the end of the decoherence time.

On the contrary, photons are well decoupled from the background, travel at the highest speed in the universe and can be precisely controlled in picosecond time resolution. Recently, photonic systems have demonstrated power in solving intractable problems like Boson sampling [4], but face challenges for practically scalable universal quantum computing solutions because it is extremely difficult for a single photon to control another deterministically. Though manipulating photonic single qubit is straightforward with linear optics [5], the path towards universal quantum computer faces a great challenge due to lack of efficient optical nonlinearity at a single-photon level. The widely used scheme with linear optics, making use of probabilistic measurement induced effective "nonlinearity", is practically not efficient for large scale implementation because it requires enormous amount of ancilla photons and computational time [6].

Opportunity with matter-photon hybrid modules

Matter-photon hybrid modules may provide an opportunity to overcome the above challenges, meeting the following two requirements: 1) a high conversion efficiency between a local matter qubit and a photonic qubit; and 2) two-qubit controlled gate operation with high fidelity. Many such hybrid modules can be scaled up to build a distributed quantum computer whose components are remotely located. In this distributed quantum computer, single-qubit operations are performed with photons through linear optics, and inter-qubit gate operations are performed with local matter qubits. The matter-photon hybrid modules are the key to realize such a scalable and extendable distributed quantum computer. This is a new concept to the existing quantum computing platforms.

Beside for building distributed quantum computers, the above matter-photon hybrid modules can also be employed to connect remote quantum computers for distributed quantum computing.

This new distributed quantum computing scheme based on matter-photon hybrid modules also naturally solve the second problem of decoherence time in the existing platforms. In our envisioned scheme, all single qubit operations are performed on photonic qubits with linear optics which do not introduce decoherence. The only problem is loss, but unlike the decoherence in the existing quantum computing platforms that reduces pure quantum states into mixed states, the loss removes photons out of their computational basis. As a result, the photon loss reduces measurement efficiency (or probability), but does not affect qubit state fidelity. This is fundamentally different from the existing quantum computing platforms. For two-qubit control gate operations, we convert photonic qubits into two local matter qubits which require only enough time to perform one control gate operation. As a result, the matter qubit decoherence time does not affect the overall quantum computation steps. Therefore, in the matter-photon hybrid scheme the computation time is only limited by the photon loss, but not by the matter qubit lifetime or decoherence time.

Assessment

To realize the above proposed matter-photon hybrid modules are not obvious, because their two requirements may not be easily met simultaneously. Although single atoms (or ions, or superconducting qubits) have been demonstrated for realizing controlled gates with high fidelity, the conversion efficiency between these computing qubits and single photons is very low. Recently atomic-ensemble-based qubits have been proposed to overcome this problem [7], but technical challenges may remain foreseen.

The performance assessment of such matter-photon hybrid modules is based on how well these two requirements can be simultaneously met: conversion efficiency and gate fidelity. The conversion efficiency of the hybrid module determines the computation steps or network depth. The gate fidelity affects its error tolerance. The power of distributed quantum computing is assessed by the dimension of its Hilbert space.

Timeliness or maturity

Distributed quantum computing based on matter-photon hybrid modules requires long-term investments, as it provides a physical layer backbone to support both quantum computing and quantum networks – quantum internet. There are new research challenges need to be addressed, from fundamental physics, advanced quantum control technologies, to algorithms and error correction.

References:

[1] J. M. Chow, Quantum intranet, IET Quantum Communication **2**, 26–27, (2021).

[2] https://coldquanta.com/core-technology/hilbert/

[3] https://ionq.com

[4] H. -S. Zhong *et. al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).

[5] S. Barz, Quantum computing with photons: introduction to the circuit model, the one-way quantum computer, and the fundamental principles of photonic experiments, J. Phys. B: At. Mol. Opt. Phys. **48**, 083001 (2015).

[6] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, Linear optical quantum computing with photonic qubits, Rev. Mod. Phys. **79**, 135 (2007).

[7] E. Oh, X. Lai, J. Wen, and S. Du, "Distributed quantum computing with photons and atomic memories," Adv. Quantum Technol. 2300007 (2023).

Quantum Computing and Networking for Energy Applications

Hari P. Paudel,[a,b] Scott E. Crawford,[a,b] Yueh-Lin Lee,[a,b] Dominic Alfonso,[a] Manh Tien Nguyen,[a] Benjamin Avramidis,[a,c] Kenneth D. Jordan,[a,c] Darren Mollot,[d] Yuhua Duan[a*]

[a] National Energy Technology Laboratory, United States Department of Energy, Pittsburgh, PA 15236, USA
[b] NETL Support Contractor, 626 Cochrans Mill Road, Pittsburgh, PA 15236-0940, USA
[c] Department of Chemistry, University of Pittsburgh, Pittsburgh, PA 15219, USA
[d] Artificial Intelligence and Special Projects, Fossil Energy and Carbon Management Office, United States Department of Energy, Washington, DC 20585, USA

* yuhua.duan@netl.doe.gov

Topic: *Applications*: New kinds of DOE science applications informed by quantum capacities
Assessment of realistic quantum advantages
Tools for application performance modeling and estimating
Application-inspired benchmarks and curated libraries of instances
Computing and programming models: Models for hybrid quantum and classical computing
Algorithms: Hybrid quantum and classical algorithms

Challenge: While quantum computing is considered as a paradigm shift in our basic understanding of physical computation, effective implementation of quantum computing in energy applications also depends on progress and development in the dimensions of both quantum computing hardware and quantum computing algorithms. From the perspectives of quantum computing hardware, the availability of the number of qubits and the noise level of the qubits should be weighed, whereas from the perspective of quantum computing algorithms, error tolerance capability in the algorithm and gain of speed-up relative to classical computing should be considered. In addition, current quantum processing devices and quantum computing algorithms may also require pre- and post-processing by classical computers for its basic operation within realistic architecture. A gap exists between the capabilities of current quantum information science (QIS) stakeholders and energy sector needs. Quantum networking is another rapidly growing direction of QIS. For safe and secure utilization of energy, a number of challenges exist where quantum networking could find opportunities to deliver improved and environmentally friendly technologies. An emerging area of quantum networking research is the application of computational techniques for quantum networking algorithm and protocol discovery for enhanced QIS performance. Performances of quantum network scenarios can be evaluated through quantum network simulations to address possible implementation challenges in scaling up quantum networks (such as bottlenecks or the dynamics of the network interacting with systems that consist of sending and receiving circuits). Further challenges exist to seamlessly interconnect quantum sensing device with quantum networks for secure transmission of data thru quantum communication channels and finally analyze transmitted data using quantum computing.

Opportunity: Fossil energy remains a major component of the global energy portfolio. Improving and enhancing the performance of existing technologies and equipment is necessary to circumvent the growing challenges of both developing sufficient supplies of energy to meet consumer needs while reducing greenhouse gas emission. Quantum computation is expected to overcome some of the specific near- to intermediate-term challenges related to fossil energy by leveraging (i) quantum many body computation of material properties to an unprecedented level of accuracy at a lower cost, (ii) a better understating of downstream, midstream and upstream fluid flows, and (iii) quantum-enhanced machine learning and artificial intelligence for (a) developing optimum materials for carbon capture technology that outperform existing materials, and (b) studying deep geological data for oil, gas, and mineral explorations. Meanwhile, quantum networking simulations for energy-related applications might be deployed in the design phase to evaluate the merits of various quantum network protocols and architectures, or for assessing real-time performances and troubleshooting of the built network. Quantum networking and communication protocols will play a significant role in the energy sector: innovations such as deployment of electric vehicle charging stations, development of smart grids, and expansion of microgrids will require sophisticated sensor networks to record data for optimized efficiency, which creates new vulnerabilities for cyberattacks and other disruptions. Quantum protocols such as quantum key distribution (QKD) will be essential for ensuring secure collection and transmission of data while protecting consumer privacy and avoiding disruptions to the energy supply from outside attacks. Similarly, nuclear plants would benefit greatly from quantum communication protocols both in external networking/data exchange and securing communication and sensor readouts from remote sensors throughout the plant. Satellite monitoring of greenhouse gas emission and/or gas leaks also relies upon sophisticated sensor networks that must be secured, an area that will benefit from ground-to-satellite QKD development. Continued demonstrations of quantum communication in harsh environments and the development of mobile quantum networks will be invaluable for the creation of market-ready quantum solutions for the energy sector.

Assessment: Quantum computers have already been used to model chemical reactions, and as this technology continues to develop it will have transformative implications for material design and discovery. The ability to model large systems and rapidly screen material properties will significantly benefit the energy sector, where materials are continually being sought to improve greenhouse gas capture, catalyze new reactions, and selectively detect analytes of interest in harsh conditions. Significant work is needed to improve the performance of existing quantum computer designs while lowering production and operating costs and increasing accessibility. The developments of quantum networking simulations, like what has been done for the classical internet, will be further enhanced as the quantum technologies get maturity. This enhancement will enable more contributors to develop open-source simulation platforms. Finally, the continued maturation and commercialization of quantum technologies will naturally accelerate innovation and advancement in the field.

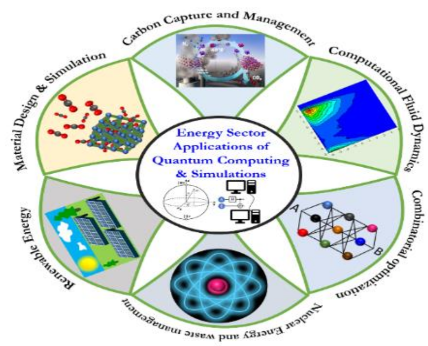

Timeliness or maturity: The energy sector will likely be among the first beneficiaries of the impending "quantum revolution," as emerging QIS enhanced technologies may be applied to ensure the safe, secure, and efficient use of energy resources. Quantum chemistry is expected to be one of the early areas that can benefit from quantum computing. Quantum networking may have the most benefits for the energy sector in the short term, given its relative maturity and the number of commercial players in the market. The ever-increasing amount of data that must be collected and disseminated within the energy sector creates the potential for vulnerability to hackers and other outside attacks, necessitating secure networks and communication protocols. In the longer term, networks of quantum sensors will lead to significant increases in sensing performance for applications such as emission monitoring, oil and gas discovery, pipeline integrity, and nuclear powerplant security. The fiber optic network may be used to test and validate new quantum technologies on an actual quantum network, a crucial step for advancement.

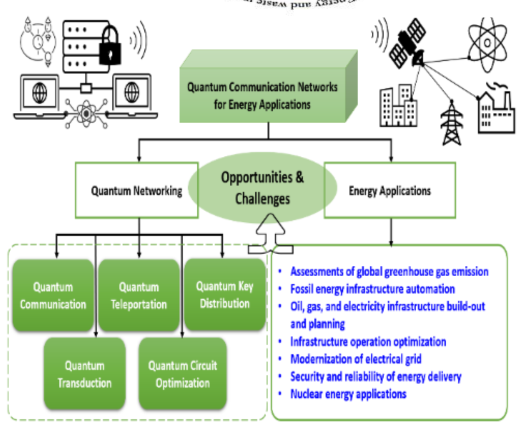

References: (1) H. P. Paudel, M. Syamlal, Y.-L. Lee, S. E. Crawford, R. A. Shugayev, P. Lu, P. R. Ohodnicki, D. Mollot, Y. Duan, "*Quantum computing and simulations for energy applications: review and perspective*", **ACS Engineering Au**, **2**(2022)151-196. (2) H. P. Paudel, S. E. Crawford, Y.-L. Lee, M. Leuenberger, R. A. Shugayev, M. Syamlal, P. R. Ohodnicki, P. Lu, D. Mollot, Y. Duan, "*Quantum Communication Networks for Energy Applications: Review and Perspective*", **Adv. Quantum Technol.** (2023) under review. (3) S. E. Crawford, R. A. Shugayev, H. P. Paudel, P. Lu, M. Syamlal, P. R. Ohodnicki, B. Chorpening, R. Gentry, Y. Duan, "*Quantum Sensor for energy applications: review and perspective*", **Adv. Quantum Technol. 4**(2021)2100049. (4) S. Bush, Y. Duan, B. Gilbert, A. Hussy, D. Miller, R. Pooser, M. Syamlal, **FE Quantum Information Science and Technology (QIST) Workshop Report**, DOE-NETL, 2020. doi: 10.2172/1639026. (5) M. T. Nguyen, Y.-L. Lee, D. Alfonso, Q. Shao, Y. Duan, "*Description of Reaction and Vibrational Energetics of CO_2-NH_3 Interaction Using Quantum Computing Algorithms*", **AVS Quantum Science 5**(2023)013801. (6) B. Avramidis, H. P. Paudel, D. Alfonso, K. D. Jordan, Y. Duan, *"Ground State Property Calculations of LiH_y Complexes using IBM Qiskit's Quantum Simulator"*, **AVS Quantum Science** (2023) to be submitted.

Integrating Application, Model, Algorithm, Compilation, and Error Correction Quantum Chasms With ~~QASM~~ Type Theory

Eugene Dumitrescu[1, *]

[1] *Quantum Information Science Section, Oak Ridge National Laboratory*

Challenge— *Efficiently* and *correctly* programming a quantum computer is difficult in part because, by connecting high-level algorithmic constructs with a physical-level hardware realization, a complete quantum computation traverses wide chasms between vastly different hierarchical computational scales. As the technical machinery to express even a simple quantum computation across *all* scales is lacking, we propose type theory as a central ingredient in efficient and correct quantum programs. Different relevant (as well as marginal and irrelevant) types, associated with instructions operating upon them, should come into focus as one re-scales along the physical- to algorithmic-levels in such a consistent hierarchical type theory. Likewise, an understanding all hierarchical levels and their properties is required for a true end-to-end compiler.

For any type to be useful a working standard operating library is also required. Standard libraries defining arithmetic operations acting on quantum types (qint, qfloat, and qetc defined later), as well as time-evolution protocols, and standard oracular gadgetry are required to construct high-level programs. Some type properties are transferable across levels, while others are not. To unify this, a core representation must exist which expresses salient attributes at each hierarchical level. While some software has been dedicated to investigating this interplay of types and data-structures [1, 2], more are needed to deliver complete type theory which is fully operational and correct (whether by proof or by unit-test). While library constructions will require great efforts, standardizing types will dramatically reduce the effort to maintain codes as well to develop a myriad of new languages, compilers, and programs.

Opportunity— To highlight and organize type theory's implications across the spectrum we add a new attribute to quantum types. That is, define a *color* quantum attribute number $\mu \in \{r, g, b\}$ denoting the quantum computing types' and instructions' levels within the physical-to-algorithmic computational hierarchy. The largest-scale, so called infrared-red (IR), scale logical operational modalities are red (r), error correction coarse-grains through the intermediary green (g) spectrum, and the microscopic physical-length-scale noisy ultra-violet (UV) types and instructions are designated blue (b). Analog (sometimes called noisy intermediate scale quantum) experiments reside in a limited near-term computational modality where this color hierarchy collapses ($\mu = r = g = b$) which indicates that the types and instruction sets used across the hardware and the logical levels are identical. As error correction first emerges there will be, for the first time, a distinction between the physical and logical types. As outlined below, we argue that realizing universal quantum computations will be composed of types and compilers throughout the spectrum.

Assessment— At the different layers of a quantum computation:

(r) Science Oriented Languages— From a *domain scientist's* perspective high-level quantum type theory is the crucial tool to formulate, compile, and compute specific problems of interest. The working domain scientist already implicitly has a good grasp of quantum types since they previously utilized classical instantiating (and approximations) of fundamental quantum types in modeling physical phenomena. Here one utilized the familiar and expressive classical types (bits, bytes, ints, floats, chars, classes, etc) defined within the context of a given programming language to models properties of, for example, quantum fields.

Already at this level, a need for fundamental quantum types for expressive and operational purposes emerges. Fundamental quantum computing types should include qubit, fermion, qumode, ancilla, interferometer, (weak) measurement, and classical groups (and eventually beyond). Such quantum types are needed to both express generic as well as design optimal quantum algorithms. In a Feynman-like approach, simulating fermions with a fermionic quantum computer [3] (equipped with type f_r) is the most straightforward path for the high-level chemist. However due to the difficulties of simulating fermions, as well as the needs for other types, one often begins with other types such as q_r, b_r, a_r for qubit (represented by Pauli-algebra), boson (canonical commutation relations), and anyon (defined by the parent quantum field theory's topological datum [4, 5]) respectively.

(g) Error Correction— Along with lowering the noise floor, error correction is the substrate mapping, distributing, and enabling a manifestly digitized quantum communication and computation. The latter component is of special interest because it maps types and instructions between the bare and encoded layers. For example, superconducting and atomic qubits have both been used to mimic small scale quantum error correcting codes. Hence we have a first example of two experiments which are fundamentally distinct at the $\mu = b$ physical-level but are logically equivalent in operation (up to error rates) using the encoded g logical types. Note that within this framework, many logical

* edq@ornl.gov

concatenations could abstractly be encapsulated at the emergent $\mu = g$ level. For example, one may envision first using a $[[n, k, d]]$ encoding, mapping the n physical to the k logical qubits $q_b \rightarrow q_{g_1}$. We could further concatenate codes to re-encode qubits, with e.g., a $[[n' = k, k', d']]$ encoding: $q_{g_1} \rightarrow q_{g_2}$. By defining a code with manifestly fermionic characteristics (other characteristics would accompany other types of emergent fields) we alternatively have an encoding map: $q_{g_1} \rightarrow f_{g_2}$ [6]. One can continue in this fashion further encoding the fermions into a higher type $f_{g_2} \rightarrow q_{g_3}$, or alternatively as might be the case for chemistry applications, just accept this type as sufficiently encoded to be regarded, and compiled to, as a high-level logical type: $f_{g_2} \equiv f_r$. Just like at the logical level, standard libraries are required for error correction.

(b) Microscopic Device Description— At the UV, physical length scale, a device implements "digital"—i.e. prespecified and hopefully well characterized analog instructions—dynamics to realize (quasi-)isolated Schrodinger-equation dynamics on a qubit's physical Hilbert space. These physical operations form the basis of either error correcting instructions or analog operations. Perhaps due to the emergence of superconducting[7, 8], atomic[9], and photonic qubits [10] (defining discrete *qudit* or continuous variable *qumode* types) a large amount of recent attention has been dedicated to formalizing *qubit* types which are various instances of q_b. Similar to how a qubit type is implicitly and then later explicitly defined by their $SU(2)$ operations and generating local Pauli algebra, a qumode type with certain instruction sets (linear optics and two-mode squeezers) could also concisely and formally be defined with data structure being the group $SU(1,1)$ [11]. Other types can and should also be defined by a succinct mathematical representation. For example, qubits are defined by space and time coordinates as well as algebraic computational operations. Here, at the level of $SU(2)$, the Euler, ZYZ Cartan, and U_3 decompositions are all synonymous (although they may be generalized in different ways). In addition, there is a further need to generically use the fundamental types as units building *compound* quantum types such as a qreg (a qubit array).

Timeliness and maturity— Using the color type attribute we can now organize and connect programs along a common (color) axis. For example, the popular pythonic framework Qiskit compiles a user's programs, with types at this bottom b-level, into a json dictionary containing a valid OpenQASM2.0 [7] compiled program. OpenQASM3.0 [12] tentatively makes a step towards our definition by separating along coherent (purely quantum types) and longer, incoherent, timescales (with low-level classical types). Such considerations are required to understand the b level and its interactions with g and r. The quantum intermediate representation (QIR), and prior works, have introduced initial LLVM compatible quantum compiler tool chains [13, 14]. However the current QIR lives close to the physical layer and does not couple to the (r, g) degrees of freedom. As with all three levels, the compilation framework should abide by principles of, or systematically reflect, linear logic and relevant properties of quantum information (no cloning, reversibility for pure systems, or contractiveness of dissipative maps). In this way, proper type theory aids facilitates the extermination quantum bugs from our programs [15]. Of course, this is in addition to type theory defining, informing, and creating the structure to operationally translate quantum computations across scales.

---

[1] Jarrod R. McClean et al. OpenFermion: The electronic structure package for quantum computers. *Quantum Science and Technology*, 5(3), 10 2020.

[2] Tiffany M. Mintz et al. QCOR: A language extension specification for the heterogeneous quantum-classical model of computation. *ACM Journal on Emerging Technologies in Computing Systems*, 16(2):1–17, 4 2020.

[3] Sergey B. Bravyi and Alexei Yu Kitaev. Fermionic quantum computation. *Annals of Physics*, 298(1):210–226, 5 2002.

[4] Chetan Nayak, Steven H. Simon, Ady Stern, Michael Freedman, and Sankar Das Sarma. Non-Abelian anyons and topological quantum computation. *Reviews of Modern Physics*, 80(3):1083–1159, 9 2008.

[5] Emil Génetay Johansen and Tapio Simula. Fibonacci anyons versus majorana fermions: A monte carlo approach to the compilation of braid circuits in $SU(2)_k$ anyon models. *PRX Quantum*, 2:010334, 3 2021.

[6] Andrew J. Landahl and Benjamin C. A. Morrison. Logical Majorana fermions for fault-tolerant quantum simulation. 2021.

[7] Andrew W. Cross, Lev S. Bishop, John A. Smolin, and Jay M. Gambetta. Open quantum assembly language. 7 2017.

[8] Robert S. Smith, Michael J. Curtis, and William J. Zeng. A Practical Quantum Instruction Set Architecture. 8 2016.

[9] Ebadi et al. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595:227–232, 7 2021.

[10] Sara Bartolucci et al. Fusion-based quantum computation. *Nature Communications*, 14(1), 1 2023.

[11] Z. Shaterzadeh-Yazdi, P. S. Turner, and B. C. Sanders. SU(1,1) symmetry of multimode squeezed states. *Journal of Physics A: Mathematical and Theoretical*, 41(5):055309, 2 2008.

[12] Andrew W. Cross et al. OpenQASM 3: A Broader and Deeper Quantum Assembly Language. *ACM Transactions on Quantum Computing*, 3(3):1–50, 4 2022.

[13] A.J. McCaskey, E.F. Dumitrescu, D. Liakh, M. Chen, W. Feng, and T.S. Humble. A language and hardware independent approach to quantum–classical computing. *SoftwareX*, 7:245–254, 1 2018.

[14] Thien Nguyen and Alexander McCaskey. Retargetable Optimizing Compilers for Quantum Accelerators via a Multilevel Intermediate Representation. *IEEE Micro*, 42(5):17–33, 9 2022.

[15] Junjie Luo and Jianjun Zhao. Formalization of Quantum Intermediate Representations for Code Safety. 3 2023.

Title: Physics-informed generalizable quantum software co-design

Authors: Avik Dutt (University of Maryland), Manish Kumar Singh (memQ), Sean Sullivan (memQ)

Topic: Quantum computing (QC) models, error mitigation, codesign and integration for quantum computing and networking, compilation.

Challenge: There has been remarkable progress in both hardware and software for QCs based on multiple models - quantum annealers (QA), gate-based digital QCs, measurement-based photonic QCs, and analog quantum simulators (AQS). Quantum annealing à la DWave has moved from questionable performance to demonstrable speedups over classical brute–force optimization and classical simulated annealing. Similarly, analog quantum simulators (AQS) have proliferated [1,2]. Quantum software is quite well developed individually for the gate-based models, for quantum annealers, for measurement-based photonic quantum computers, and for AQS. Building software stacks that can harness the benefits of each of these platforms, interface between them in a distributed fashion, while also connecting to classical supercomputers, is in its infancy. For such a distributed platform would require quantum networking software to be developed that can model and control the hardware of the quantum network aspects. Hence we identify two major challenges from a quantum software stack perspective: (1) cross-platform and cross-model quantum computing software interfaced with powerful classical computers, that can efficiently transpile and distribute parts of QC circuits that are best suited for each respective hardware platform, and (2) quantum networking software – to enable reliable modeling and control of the networking hardware (quantum repeaters, QuICs [3] and components, i.e., source, detectors, etc.) – is less developed/ nonexistent.

Opportunity: We discuss the opportunities afforded by these challenges in two parts:
(1) The co-design of software alongside deep understanding of the underlying physics of the hardware, and co-design of QC software stacks across not only across different qubit platforms but across different QC models presents a great opportunity. However, this is a formidable challenge to overcome due to the very different modes of operation of each QC mode: qubits vs analog simulators vs quantum annealers. Such software would still be limited to small modular quantum computing nodes.

(2) One common challenge between the leading platforms for quantum computing, when it comes to further scaling up useful QCs, is the networking problem. While promising systems of 100s qubits have been demonstrated, they are far from the points where these computers become useful. The path to a large number of qubits is either not clear or presents immense challenges. It is therefore important that we also look at technologies that are natively scalable when it comes to quantum networking and have the potential to become a distributed quantum computing platform connecting modular QC nodes. Software for quantum networking is much less developed than for quantum computing, where a number of open source alternatives exist.

An example is solid-state qubit based systems [3] such as NV centers, rare-earths [4], SiC defects etc - being actively pursued in private (Amazon, Psiquantum, memQ, HRL etc.) as well as academic research. In order to speed-up the development of these platforms, we need to simultaneously develop accurate software device models to guide the design of physical systems, and inform the required performance metrics of individual components and entire systems. This is extremely important as even “deterministic” quantum dot systems effectively become probabilistic due to the inefficiencies introduced by coupling, loss and other imperfections. Of equal importance is the development of protocols and techniques to address hardware for mitigating qubit errors through optimized control sequences, keeping in mind device/system parameters that are achievable in the near term (~5 years).

Assessment: The following key, tangible development provide a good way to assess the solutions and their success:

- Co-designed solutions with hardware developers.
- Development of optimized control and calibration of photonic quantum computing hardware in simulation. Availability of open-source software packages that fill the identified gap (Figure 1).
- Software-driven /data-driven identification of system parameters with the largest impact on computing and quantum networking performance, and their feed-forward into respective technology platforms.
- Adoption by hardware & software technology developers in academia or industry.

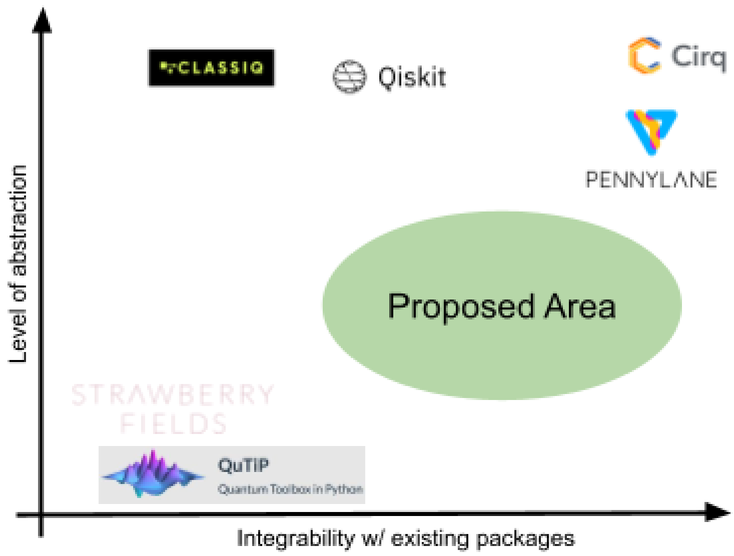

Figure 1. There is a gap between existing low-level quantum optics device simulators (QuTiP, QuantumOptics.jl, Strawberry Fields) and abstract qubit circuit simulators, optimizers and transpilers (Qiskit, Pennylane, Classiq)

Timeliness or maturity: As mentioned before, software for quantum networking is much less developed than for quantum computing, where many projects exist such as IBM Qiskit, Circ, Amazon Braket, Pennylane to name a few on the algorithmic side; and QuTiP, QuantumOptics.jl on the low-level physical hardware side. Both academic and industry research scientists are poised to harness a mid-level quantum software stack that bridges the gap between the high-level, generic software and low-level hardware-specific software packages as shown in Figure 1. Since a majority of the codebase of existing QC software is open-source, the time is ripe and sufficient maturity is available to build this mid-level software stack in a hierarchical fashion. Quantum networking would be intimately connected with this co-designed hardware-software stack, and hence such a mid-level modeling, simulation and control package would directly feed into the development of a quantum networking software platform. Other reasons for the timeliness of such developments are that an understanding of low-level physical hardware technologies [4] has only matured in the last few years, and so has the ability to reliably predict their behavior. This is essential for the hierarchical software platform that is necessary for distributed QC and quantum networking. Quantum photonic technologies stand to especially gain from this effort as they can be used both for measurement-based quantum computation and are the preferred medium for quantum networking.

References

1. J. I. Cirac and P. Zoller, "Goals and opportunities in quantum simulation," Nat. Phys. **8**, 264–266 (2012).
2. S. Ebadi,et al., "Quantum phases of matter on a 256-atom programmable quantum simulator," Nature **595**, 227–232 (2021).
3. D. Awschalom,et al., "Development of Quantum Interconnects (QuICs) for Next-Generation Information Technologies," PRX Quantum **2**, 017002 (2021).
4. A. M. Dibos, et al., "Purcell Enhancement of Erbium Ions in TiO2 on Silicon Nanocavities," Nano Lett. **22**, 6530–6536 (2022).

Daisy-Chaining Annealing to Hamiltonian Simulation
for Ground State and Dynamics Calculations

Stephan Eidenbenz*, John Golden
Los Alamos National Laboratory
*eidenben@lanl.gov

Topic: applications, models, algorithms

Challenge:
Hamiltonian simulation is the key algorithmic building block of most quantum computing applications. At its heart, Hamiltonian simulation solves the Schroedinger equation and allows us to simulate dynamics of a physical system defined by a time-dependent or time-independent Hamiltonian. In fact, simulating complicated many-body Hamiltonians was the original use case for quantum computers as laid out by Richard Feynman in 1982 [1]. A variety of algorithmic methods to implement Hamiltonian simulations have been developed for the circuit model of quantum computing over the past decades including Trotterization, Linear Combination of Unitaries (LCU), and Quantum Signal Processing (LCU). In total, Hamiltonian simulation has seen by far the most effort in terms of algorithmic development of any computational quantum method. Hamiltonian simulation is also the only algorithmic problem – other than factoring, where quantum approaches are generally expected to be able unquestionably obtain an exponential speed-up over their classical counterparts.
Envisioned use cases for Hamiltonian simulation include, to name only a select subset, (i) studying superconductivity through the Fermi-Hubbard model [2], (ii) properties of magnetism through spin chains [3], and (iii) scattering processes in quantum chromodynamics [4]. All of these use cases seek to advance our fundamental understanding of physics by simulating through time, in other words *understanding the dynamics* of the underlying physics processes (or simplified models thereof). Most use cases assume that we start in a standard reference point, which is usually the ground state of the Hamiltonian. *They key challenge, however, is that these ground states are unknown in most cases and certainly for the three specific use cases mentioned.* While some algorithms exist for preparing the ground state on a gate-based quantum computer, they generally require either the prior knowledge of a state with significant overlap with the ground state [5], or significant quantum and classical computational resources to variationally discover the ground state [6]. However, there is a much more systematic and elegant method.

Opportunity:
Quantum annealing is a perfect candidate for finding ground states of a Hamiltonian. Quantum annealers start in the known (and easily prepared) ground state of a simple Hamiltonian and then slowly (ie adiabatically) phase in the target Hamiltonian resulting in the quantum ground state [7]. Quantum computers built specifically for the purpose of annealing are much simpler than general gate-based machines, and thus obtain ground states much quicker and with greater accuracy. *We propose that a ground state could be prepared by a quantum annealer and then fed directly into a gate-level quantum computer.* This would require transferring a quantum state either through a direct physical connection from the

quantum annealer to the gate-based quantum computer, or potentially via a quantum networking channel (without breaking the no-cloning theorem).

Assessment:
If successful, this approach would be a clear *path towards a practical quantum advantage for scientific applications.* Daisy-chaining annealing to Hamiltonian simulation could show at a stage before error-correction becomes dominant that analog quantum computing through annealing and gate-level quantum devices could benefit from each other. Particularly, this approach could also provide a challenge for quantum networking. Further exploring this approach could seed new approaches to developing hybrid, or cross-platform and cross-compute model quantum algorithms.

Timeliness or maturity:
NISQ devices (such as IBMQ devices or Quantinuums H-1) continue to improve their quality, for instance in terms of quantum volume and other benchmarks, whereas quantum annealers, such as D-Wave's devices continue to improve their technology. Both technologies have started to grow at a rate similar to Moore's law.

References:

[1] Feynman, R.P. "Simulating physics with computers." *Int J Theor Phys* **21**, 467–488 (1982).

[2] Arute, Frank, et al. "Observation of separated dynamics of charge and spin in the fermi-hubbard model." *arXiv preprint arXiv:2010.07965* (2020).

[3] Crippa, Luca, et al. "Simulating static and dynamic properties of magnetic molecules with prototype quantum computers." *Magnetochemistry* 7.8 (2021): 117.

[4] Davoudi, Christian W. Bauer, et al. "Quantum simulation for high energy physics." *arXiv preprint arXiv:2204.03381* (2022).

[5] Lin, Lin, and Yu Tong. "Near-optimal ground state preparation." *Quantum* 4 (2020): 372.

[6] Tilly, Jules, et al. "The variational quantum eigensolver: a review of methods and best practices." *Physics Reports* 986 (2022): 1-128.

[7] Farhi, Edward, et al. "Quantum computation by adiabatic evolution." *arXiv preprint quant-ph/0001106* (2000).

Towards Distributed Quantum Computing:
Developing a hybrid quantum interconnect stack

Dirk Englund (MIT/BNL, englund@mit.edu), Stefan Krastanov (MIT, stefankr@mit.edu), Matthew Eichenfield (U. of Arizona, eichenfield@arizona.edu), Chuck Choi (MIT, choihr@mit.edu), Sebastian Will (Columbia University/BNL, sw3151@columbia.edu), Dimitrios Katramatos (BNL, dkat@bnl.gov), Joanna Zajac (BNL, jzajac@bnl.gov), Layla Hormozi (BNL, hormozi@bnl.gov), Gabriella Carini (BNL, carini@bnl.gov), Tim Stavenger (PNNL, timothy.stavenger@pnnl.gov), Hong Tang (Yale University, hong.tang@yale.edu), Kathy-Anne Soderberg (AFRL, kathy-anne.soderberg@us.af.mil), Eden Figueroa (Stony Brook University/BNL, eden.figueroa@stonybrook.edu).

Topic: Quantum Computing and Networking Models

1. Challenge: Scaling Quantum Processors.

The deployment of large, scalable quantum computing systems is one of the most coveted targets in Quantum Information Science and Technology. The development of current Noisy Intermediate Scale Quantum (NISQ) computing systems has focused on engineering large amounts of addressable physical qubits that exist in the same space (such as the 433-qubit IBM Osprey or the 256-qubit QuEra quantum processors). In this paradigm, entanglement is created locally by applying near-neighbor quantum gates between qubits in close proximity. The execution of quantum algorithms is thus achieved using the *quantum computing stack* [1], which is an abstraction defining the flow between the user-defined quantum algorithms, error correction, and the quantum computing hardware. Because the resources needed to entangle many qubits simultaneously increase significantly with the number of qubits, it is quite challenging to extend the near-neighbor paradigm beyond hundreds of qubits. A new paradigm is required to go beyond the physical limits of current quantum computing systems.

2. Opportunity: Quantum Processing Through the Quantum Internet.

Recent approaches have proposed applying the concept of non-local entanglement to the generation of pairs of entangled qubits within a quantum processor unit [2]. The clear advantage of such approaches is that, in this way, the entanglement can then be routed in photonic form, allowing for the creation of entangled pairs of qubits that are beyond the near-neighbor limit. Non-local entangled photon pairs can then be used to drive teleportation processes between different quantum processors [3]. In this new paradigm, the use of teleportation-assisted algorithms is executed through the *quantum internet stack*, which is an abstraction defining the flow between user-defined, in-network quantum algorithms, and entanglement distribution hardware [4]. The physical layer of such a hybrid quantum processing architecture will require the integration of quantum processing (QPUs), large-scale quantum memory banks, QRAM, user-defined entanglement distribution networks, efficient gate synthesis techniques, and scalable quantum-aware control systems. Developing this new Distributed Quantum Computing (DQC) stack will require a co-design approach: The development of a more mature quantum internet stack driving quantum-repeater-based

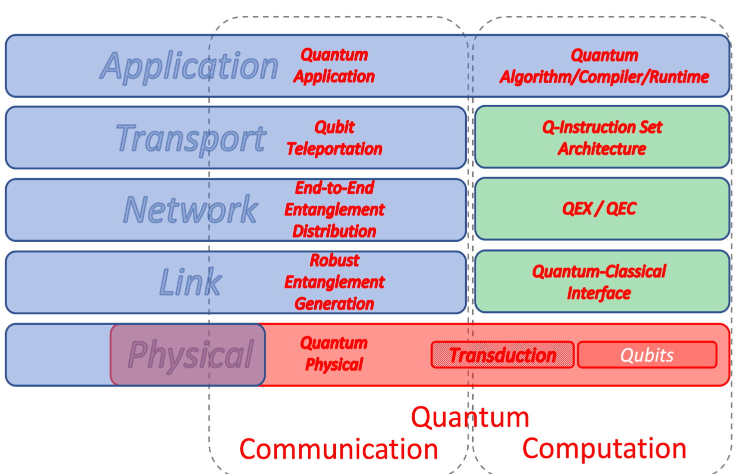

entanglement connections should be closely coordinated with the development of a more advanced quantum computing stack driving quantum computers with a large number of qubits. Reliable sharing of entanglement between QPUs will require the development of highly efficient quantum frequency conversion devices. Additionally, a co-design approach must be followed in matching quantum applications with the DQC quantum hardware at the physical layer.

Utilizing the co-designed DQC stack will further involve: (i) developing quantum algorithms with evidence of advantage in fundamental domains, such as simulation, optimization, and machine learning [5], (ii) investigating hybrid quantum and classical algorithms to capitalize on the strengths of both paradigms, for example using quantum speed-up in reinforcement learning agents [6], (iii) exploring quantum-inspired classical algorithms to enhance classical computing tasks, including convolution, cross-correlation, and equivariant transformations [7], and (iv) designing classical algorithms and software systems to simulate quantum computers and networks, including tensor network and Monte Carlo simulations [8].

3. Assessment: A Hybrid Quantum Computing/Internet Stack.

We envision the implementation of new quantum processing systems consisting of several qubits, that by design will be compatible with the principles of entanglement generation and distribution that are widely used in the QIST community [9]. These next-generation systems will allow us to characterize the success of the proposed distributed architecture by measuring its ability to efficiently execute quantum algorithms and provide private, fast services for quantum computation, communication, and sensing. This includes assessing their performance using multi-party private quantum communication, and distributed sensing through data compression. Additionally, the success of the new co-designed distributed quantum computing algorithms will be measured by benchmarking their ability to overcome the limitations of current quantum computing and networking stacks. Potential solutions will be evaluated based on their performance and resource optimization, as well as their impact on the broader field of quantum computing. Metrics for success may include the scalability of algorithms, the reduction of error rates, and improvements in qubit connectivity.

4. Maturity: Hybrid Stack Realization.

Co-designing of quantum algorithms and hardware is a timely and essential research direction that can significantly enhance the performance and resource optimization of quantum computing and networking tasks. In addition to the necessary software and algorithms, it is important to develop the engineering of the control systems, to execute a first instance of this *hybrid quantum computing/internet stack.* Critically, control mechanisms must be designed with the perspective of integrating the current quantum computing stack concepts into the wider quantum internet stack. This new *distributed quantum computing stack* will allow us to explore the user-defined selection of which pairs of qubits are entangled locally and later extracted outside of the QPUs. First demonstrations of this new paradigm of teleportation-assisted Distributed Quantum Computing will pave the way towards scaling NISQ quantum computing systems beyond their current technical limitations. This will be an important step to potentially run distributed versions of important quantum algorithms, such as Shor's and Grover's, that require thousands of simultaneously entangled qubits to be executed. This important milestone can pave the way for practical applications in various fields, such as cryptography, drug discovery, and optimization problems, unlocking the true potential of quantum computing.

References

[1] Physics Today 74, 3, 28 (2021). [2] Nature Physics 9, 345 (2013). [3] ACM SIGACT 39, 77 (2008). [4] New J. Phys. 21, 033003 (2019), [5] Phys. Rev. A 105, 022415, (2022), [6] Nature 591, (2021), [7] Phys. Rev. A 106, 032402,(2022), [8] Nature Physics 16, (2020), [9] PRL 125, 260502 (2020).

RapidQE^2: Towards Rapid Error Estimation on NISQ Machines

Bo Fang[1]*, Bo Peng[1], Ang Li[1], Wang Meng[1,2] , Prashant Nair[2], Qiang Guan[3], Weiwen Jiang[4]
[1]Pacific Northwest National Laboratory [2]University of British Columbia [3]Kent State University [4]George Mason University
bo.fang@pnnl.gov

Call: Position Paper for Workshop on Quantum Computing and Networking
Topic: quantum error correction and mitigation, compilation

Challenge

NISQ devices provide a valuable near-term platform for testing and exploring the potential of quantum computing; however, their ability to bridge the knowledge and capability gaps required for fault-tolerant quantum computing is limited by various obstacles, such as noise and decoherence. In particular, noise, which refers to the susceptibility of qubits and quantum operations to environmental interference like temperature and electromagnetic fields, poses a significant challenge for NISQ devices. To effectively utilize NISQ machines, it is essential to understand the quality and correctness of quantum algorithm outcomes, which are primarily represented through gate-based quantum circuits.

Although major quantum machine vendors such as IBM, Rigetti, and IonQ prioritize noise characterization of each physical quantum machine via various techniques such as randomized benchmarking [1], a fundamental gap exists in extending machine-level error characteristics to specific circuit executions. Quantum noise simulation, among other methods and strategies, is a promising technique that can help address this gap. It offers benefits such as the integration of state inspection with specific noise models, the flexibility to explore different architectures, basis gates, and topologies for noise mitigation, and the capability to investigate the impact of noise on deep circuits beyond the limitations of NISQ machines. However, quantum noise simulation faces significant challenges.

- It requires a density matrix representation of the quantum state, resulting in an enormous memory footprint on the classical system. For instance, the density matrix scales as 4^n in size, where n denotes the number of qubits, and simulating a 30-qubit system consumes over 16 Exabytes of memory.
- Running quantum state simulations embedded with noise model(s) on large classical memory can be computationally expensive for any meaningful quantum circuits, particularly when considering the time resources required for complex simulations.
- The identified noise models, although meaningful mathematically, are still approximations of real-world noise channels. This results in a potentially biased understanding that cannot be completely resolved. Consequently, the characteristics of single or combined noise models may not exactly match the error characteristics of an actual quantum machine [2], making it challenging to achieve accurate simulations.

Opportunity

This paper introduces RapidQE^2, a comprehensive set of approaches aimed at fostering a new direction in the rapid estimation of quantum error probabilities for quantum circuits on NISQ machines. While exact noise simulation techniques can yield highly accurate results, RapidQE^2 seeks to balance accuracy and efficiency in noise characterization, enabling a broader understanding of quantum circuits' behavior under machine noise. This innovative approach aims to inspire collective efforts in developing noise mitigation strategies and propel the quantum computing field forward.

The potential impact of our rapid noise estimation approaches is far-reaching. By providing relatively accurate estimates of the noise impact, we can gain a deeper understanding of the intricate relationships between qubits and quantum circuits, leading to more efficient resource allocation and improved computation accuracy. Moreover, when integrated with quantum circuit transpilation/compilation, a rapid estimation can inform the optimization process for more error-resilient qubit mapping and swapping, enabling the creation of more noise-resilient circuits tailored to specific quantum machines. This empowers researchers and engineers to optimize quantum circuits under a given set of noise conditions, accelerating advancements in quantum computing. Furthermore, our approaches facilitate the identification of circuit components that are more susceptible to noise, promoting more careful design and error correction.

Drawing on the potential impact of rapid noise estimation approaches, it is instructive to examine the connections and similarities between classical and quantum error resilience characterization, as well as the specific directions that RapidQE^2 will explore. In classical fault injection, statistical techniques are used to inject faults into the program state, and the program is executed multiple times to determine the impact of the faults on the program. Quantum noise simulations, particularly when considering the full range of quantum state evolution, are computationally expensive [9]. The goal of RapidQE^2 is to reduce this computational burden by employing techniques akin to Architecture Vulnerability Factor (AVF) analysis and related Program Vulnerability Factor (PVF) and extended Program Vulnerability Factor (ePVF) analyses [3,4,5], which model inherent error masking and error propagation in the computation system's states.

RapidQE^2 has three main directions of exploration.

- Leveraging the computational graph of a transpiled quantum circuit and the error rates exposed by quantum machine vendors to compute the *likelihood of correct* states being affected by noise per operator. Prior work [7,8] demonstrates that the final quantum states can be predicted with a probabilistic model and more recent machine learning models (e.g. GNN + Transformer).
- Identifying the error-prone portion of a quantum circuit by examining the patterns of the sequences of operators applied on the quantum state and generating relative ranking over the patterns to represent how sensitive the portion of a circuit is to the noise [5].
- Investigating the error resilience of different computational basis states within quantum states when various initial states are prepared, as some basis states may exhibit greater resilience to errors introduced by quantum operations [10]. By focusing on the more susceptible states and their interactions with quantum operations, the rapid quantum error estimation process could be made more efficient without the need to exhaustively analyze all possible states. Besides, making use of the error-resilient bases also helps guide the design of more efficient and noise-resilient quantum circuits and more targeted noise mitigation strategies.

Assessment

The Quantum Vulnerability Factor (QVF) is a visionary metric that holds the potential to revolutionize the approximation of correct quantum states for a given quantum circuit on a NISQ machine. Employing this metric, RapidQE^2's success can be evaluated by correlating the QVF of a circuit on a quantum machine with the final quantum state fidelity. By identifying the more error-prone portions of a circuit, RapidQE^2 paves the way for groundbreaking evaluations using actual or simulated noise impact, revealing sensitive deviations in fidelity. This innovative approach offers valuable insights into quantum circuit behavior under diverse noise conditions, propelling the development of effective noise mitigation strategies.

Timeliness or maturity

As NISQ machines enter a transformative phase with an increasing number of qubits, error correction and mitigation techniques must adapt by concentrating on a reasonable scope of the system for efficient protection. Outperforming quantum tomography [8] in generating practical error characteristics, randomized benchmarking has emerged as a vital tool in this new era. It provides specific error rates on qubits, gates, and measurements, enabling the inference of overall quantum state fidelity. Embracing this timely and maturing technology, we are poised to unlock the full potential of quantum computing and reshape the future of the field.

References

[1] K. Emanuel, et al. "Randomized benchmarking of quantum gates." Phys. Rev. A 77.1, 2008 [2] https://ionq.com/docs/get-started-with-hardware-noise-model-simulation [3] S. Mukherjee et al, "A systematic methodology to compute the architectural vulnerability factors for a high-performance microprocessor," in Proceedings MICRO-36., pp. 29–40, 2003. [4] V. Sridharan and D. Kaeli, "Using hardware vulnerability factors to enhance avf analysis," HPCA 2010. [5] **B. Fang**, et al., "EPVF: An enhanced program vulnerability factor methodology for cross-layer resilience analysis," DSN 2016. [6] Z Yu, Y Li, "Analysis of Error Propagation in Quantum Computers", arXiv:2209.01699, 2022. [7] H. Wang, … **W. Jiang** …et al. "QuEst: Graph Transformer for Quantum Circuit Reliability Estimation." *arXiv:2210.16724,* 2022. [8] D'Ariano, Quantum Tomography. 2008. [9] **A. Li**, et al., "Density Matrix Quantum Circuit Simulation via the BSP Machine on Modern GPU Clusters", pp. 1-15, SC 2020. [10] J. C. Aulicino, T. Keen, **B. Peng**, "State preparation and evolution in quantum computing: A perspective from hamiltonian moments." Int. J. Quantum Chem. 122(5):e26853, 2022.

The nonequilibrium many-body problem: An application where quantum computers will show advantage over classical computers

James K. Freericks,[1] Ryan Bennink*,[2] and Yan Wang†[2]

[1] *Department of Physics, Georgetown University, Washington, DC,* james.freericks@georgetown.edu
[2] *Oak Ridge National Laboratory, Oak Ridge, TN,* *benninkrs@ornl.gov, †wangy2@ornl.gov

Topics: applications, algorithms, and error correction and mitigation

Challenge. Quantum computing has been long seeking a science application where quantum computers offer a significant advantage over classical computers. Usually, this advantage is sought in terms of the compute time, as determined by the number of steps required to run the algorithm. But, many domain-science and mathematical applications in classical high-performance computing, such as those from quantum chemistry and physics, are memory limited. We argue that quantum advantage may sooner be found in terms of *memory*, that is, by using quantum computers to solve problems whose underlying objects are too large and complex to be stored on a classical computer. This realm is easily reached in many-body physics applications, which involve quantum states whose representation is exponentially large in the number of particles (even using clever methods to simplify the wavefunctions for these states).

In these applications, while many practical classical algorithms exist for finding ground states and their energies, or for calculating the response of the system to external fields, nonequilibrium dynamics remain particularly challenging, because the dynamics involves contributions from a large number of energy eigenstates. Whether using numerically exact methods such as nonequilibrium dynamical mean-field theory or density-matrix renormalization group methods to compute time correlation functions, data computed at all previous time steps must be stored and thus simulations are usually limited to a small number of time steps before memory issues halt the system from being able to be evolved further in time. It is here where quantum computers can offer a true advantage—solving problems at the space-time scales that cannot even be attempted on classical computers. However, there are few quantum algorithms addressing nonequilibrium dynamics.

Opportunity. The so-called driven-dissipative many-body problem, where a system is driven by an external field and is also in contact with a reservoir, is a problem that is particularly well-suited for a quantum computer. As shown in Fig. 1, a driven-dissipative system is intrinsically stable [1]. The stability of such systems on quantum computers arises because the steady state is typically a mixed state that has no coherence in its density matrix. Furthermore, because the time-evolution map is contractive, any error in the previous time step will be cured in the next time step. Thus if one can run a single Trotter time step with sufficient fidelity, one can run as many Trotter steps as desired. As shown in the figure, proper post-processing can remove many of the residual errors due to decoherence during a single Trotter time step.

The open problem that needs to be resolved is to find an efficient implementation of the driven-dissipative many-body problem on a quantum computer when the system is strongly interacting. The conventional approach requires the implementation of D^2 Kraus operators in each time step, where D is the dimension of the Hilbert space; this is impractical. New approaches are needed to implement physically realistic dissipative dynamics using a practically small number of Kraus operators, or using other strategies to implement the nonequilibrium dynamics.

Assessment. A driven-dissipative system will either have a time-independent steady state, or a periodically varying steady state (as shown in the Fig. 1). Secondly, single-particle Green's functions satisfy sum rules for short relative times. This provides two direct checks for the algorithm. In addition, if the simplified time evolution is fully consistent with the rules of thermodynamics, the system will evolve to a thermal state when the field is turned off. These checks can be easily verified as being

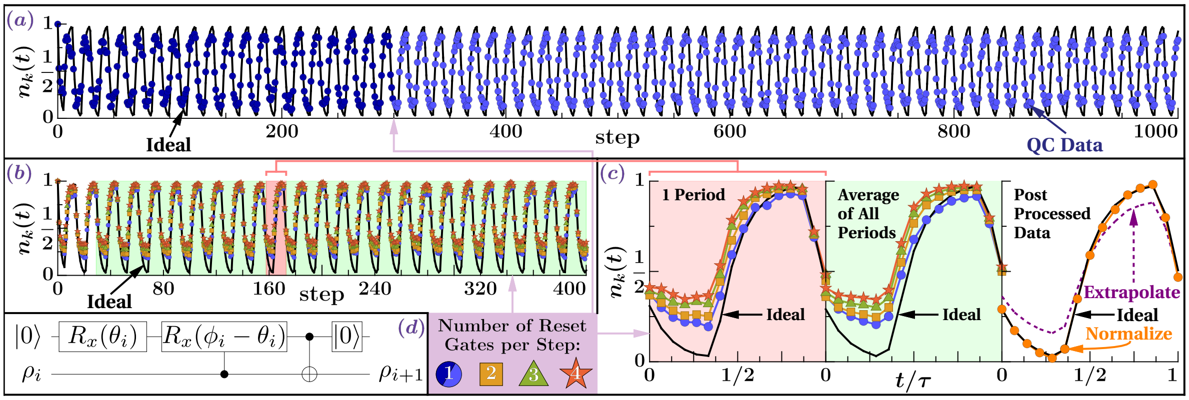

FIG. 1. Driven-dissipative dynamics of one-dimensional lattice electrons driven by a dc electric field and coupled to a reservoir. (a) Stability of the time evolution over 2000 Trotter time steps. (b) Loss of fidelity when multiple resets are applied. (c) Scaling of results with multiple time steps to the zero-reset-time limit, and further post-processing to construct the exact signal. (d) Circuit employed for each Trotter step.

correct, especially for smaller systems. The true test of assessment, however, will be to determine new phenomena. It is long believed by the community that driven-dissipative systems will evolve to non-trivial new phases of matter called nonequilibrium steady states. Finding and studying the behavior of these systems and understanding better the nature of the nonequilibrium steady state will be the ultimate assessment of the approach. Finally, any algorithm that is stable for driven-dissipative systems will also be able to examine pump-probe experiments. The DOE has significant investments in examining pump-probe experiments via their user facilities. Quantum computers can provide the means to compute theoretical predictions and theoretical verifications of these experiments.

Timeliness. Many-body physics problems have long been, and long will be, a staple of high performance computing. Although there has been much work in unitary quantum algorithms for several types of many-body physics problems, a clear quantum advantage has yet to be realized. Meanwhile, quantum algorithms leveraging dissipation—naturally occurred or engineered—remain relatively unexplored. An efficient quantum algorithm to solve the driven-dissipative many-body problem would enable significant breakthroughs in science across a number of fields including condensed matter physics, quantum materials, quantum chemistry, and high-energy physics, all of which DOE has significant investment in. While many applications await fault-tolerant architectures that may take decades to achieve the performance needed for quantum advantage, near-term quantum computers may already be sufficiently robust to be able to simulate this class of problems. This opportunity to realize the promise of quantum computing should not be missed.

[1] B. Rost, L. D. Re, N. Earnest, A. F. Kemper, B. Jones, and J. K. Freericks, Demonstrating robust simulation of driven-dissipative problems on near-term quantum computers (2021), arXiv:2108.01183 [quant-ph].

Making the Promise of Quantum Computing Practical

Jeannette M. Garcia and Antonio Mezzacapo
IBM Quantum, Almaden Research Center, San Jose, CA 95120, USA and
IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA
(Dated: May 1, 2023)

Topics: Circuit Classes, Applied Error Mitigation and Circuit Knitting

A. Challenge

Classical algorithms fail to solve some hard computational problems within practical timeframes, with the requisite accuracy for industrial use cases, or can be prohibitively expensive to execute. Therefore, identifying problems that benefit from quantum algorithms and can be solved advantageously using quantum computers is a vibrant, active area of research.

To date, evidence that quantum computing can provide computational advantages exists in three main areas with different timelines and promises for speedup:

(1) Mathematics and processing data with structure exploits a large parameter space (Hilbert space) that can be accessed by quantum computation. With the kernel method, user data can be represented and mapped to a quantum-enhanced feature space through non-linear mapping by a quantum circuit, which allows users to reveal relations in the data by a simple linear classification. Exponential speedup has been proven for a particular case of the quantum kernel approaches [1], suggesting that the applications of the quantum kernel method can be expanded to areas that benefit from machine learning approaches in which we could exploit specific data structures. Similarly speed up has been suggested for topological data analysis [2]. Linear and nonlinear differential equations are also included in this category [3][4].

(2) Simulating nature involves computing properties of quantum mechanical systems, such as an energetics, found in nature such as in chemistry, physics, biology, material sciences, etc. Electronic configurations required to describe materials grow combinatorially and often have non-classical correlations (entanglement). Quantum computing can efficiently represent quantum states with entanglement, a critical quantum mechanical property that can cost exponentially more computing resources to represent with a classical computer. Simulation of time evolution is polynomially expensive on quantum computers. Exponential speedup is not, in general, guaranteed for solving ground and thermal states properties for any system, but there is a strong belief that it will be a feature of natural systems with local interactions. Initial work on mapping molecular compounds to quantum computers involved the use of hardware-efficient circuits with variational algorithms such as variational quantum eigensolver [5].

(3) Search and optimization problems use quantum algorithms that promise a quadratic speedup such as amplitude estimationor heuristic algorithms that use a quantum Hamiltonian [6]. Even though these currently available algorithms do not offer exponential speed up for this circuit class, these approaches are commonly used in industries in which incremental speedups can still result in significant savings to users, such as in Finance.

We define quantum advantage over three variables: runtime, cost, and quality. This approximately means that a quantum computation that is faster, cheaper, better than a classical one will be advantageous.

Partnerships and scientific advances in quantum hardware, software, and theory are leveraged in real time for applications research. However, successfully evaluating instances of quantum advantage can only be achieved with thorough studies that draw fair and accurate resource comparisons to state-of-the-art classical methods; notably, classical methods are themselves continually evolving. Additionally, research requires the correct interpretation and application of enabling algorithms and software in order to glean insights with devices available today.

B. Opportunity

One of the biggest opportunities are demonstrations of quantum applications on real hardware. Error mitigation, that includes techniques such as readout error mitigation [7], zero noise extrapolation [8], and more recently, probabilistic error cancellation [9], represents a new tool that could provide utility of the hardware before fault tolerance. Quantum computational science research carried out on noisy quantum devices today leverages a suite of algorithms that improve the results and accuracy of the output from quantum computers. Error mitigation is thus commonly implemented to improve the accuracy of results from solutions using quantum hardware [10, 11]. The improved accuracy on the quantum computing outcomes can be beneficial for applications in the spaces of "simulating nature" and "mathematics and processing data with structure."

C. Assessment

We expect that some of the first quantum circuits to demonstrate the potential of the hardware and leverage error mitigation and circuit cutting techniques will be in the areas that fall under the "simulation of nature" or "mathematics and processing data with structure" categories. Simulating electronic structure problems accurately using quantum computers often requires deep quantum circuits to capture all contributions to the correlation energy, even for small molecular systems with strong correlations. On the other hand, there exist families of quantum circuits representing quantum kernels that are expected to be hard to simulate using classical computers but could be simulated on quantum computers with shallower depth compared to the electronic structure problem. A thorough comparison of runtime analysis between chemistry and machine learning problems that are stipulated to provide quantum advantage and testing on hardware will pave the path towards better understanding of where we expect to see quantum advantage first. The first efforts towards showcasing the impact of the combination of improved hardware and error mitigation techniques will be to identify areas and use cases where the hardwaresoftware synergies work better together to achieve the maximum improvement. This will involve combined efforts between theoretical investigation, device characterization and experiments with heuristics.

D. Timeline and maturity

Recent advances in hardware, error mitigation and algorithmic developments suggest that now is the time for the use of quantum computation to explore solutions to problems in quantum computational science. Specifically, we now have the ability to test circuits on quantum hardware exceeding 100 qubits and more, with device quality now reaching sub-one percent error per cycle over 100+ qubits. Moreover, recent advances in error mitigation and error suppression techniques, when coupled with circuit cutting and knitting techniques, offer researchers a suite of new tools to scale up the size of experiments that can be executed on quantum hardware. Exploring how well various error mitigation strategies perform on realistic algorithms will help narrow in on a time line for practical quantum applications. Our combined approach, if successful, will result in (1) identification of instances of quantum advantage and (2) the largest scale demonstrations to date for circuits relevant for physics, chemistry, materials science, machine learning and mathematics.

Contact email: mezzacapo@ibm.com

References

[1] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. "A rigorous and robust quantum speed-up in supervised machine learning". In: *Nature Physics* 17.9 (2021), pp. 1013–1017.

[2] Ryu Hayakawa. "Quantum algorithm for persistent Betti numbers and topological data analysis". In: *Quantum* 6 (Dec. 2022), p. 873. ISSN: 2521-327X. DOI: 10.22331/q-2022-12-07-873. URL: https://doi.org/10.22331/q-2022-12-07-873.

[3] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. "Quantum algorithm for linear systems of equations". In: *Physical review letters* 103.15 (2009), p. 150502.

[4] Jin-Peng Liu et al. "Efficient quantum algorithm for dissipative nonlinear differential equations". In: *Proceedings of the National Academy of Sciences* 118.35 (2021), e2026805118.

[5] Abhinav Kandala et al. "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets". In: *nature* 549.7671 (2017), pp. 242–246.

[6] András Gilyén, Srinivasan Arunachalam, and Nathan Wiebe. "Optimizing quantum optimization algorithms via faster quantum gradient computation". In: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2019, pp. 1425–1444.

[7] Sergey Bravyi et al. "Mitigating measurement errors in multiqubit experiments". In: *Physical Review A* 103.4 (2021), p. 042605.

[8] Abhinav Kandala et al. "Error mitigation extends the computational reach of a noisy quantum processor". In: *Nature* 567.7749 (Mar. 2019), pp. 491–495. DOI: 10.1038/s41586-019-1040-7. URL: https://doi.org/10.1038%2Fs41586-019-1040-7.

[9] Ewout van den Berg et al. *Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors*. 2022. arXiv: 2201.09866 [quant-ph].

[10] Jin Ming Koh et al. "Experimental realization of a measurement-induced entanglement phase transition on a superconducting quantum processor". In: *arXiv preprint arXiv:2203.04338* (2022).

[11] Mario Motta et al. "Quantum chemistry simulation of ground- and excited-state properties of the sulfonium cation on a superconducting quantum processor". In: *Chemical Science* 14.11 (2023), pp. 2915–2927. DOI: 10.1039/d2sc06019a. URL: https://doi.org/10.1039%2Fd2sc06019a.

Title: Quantum-Enabled Computational Science and Engineering

Authors:
Peyman Givi, peg10@pitt.edu
University of Pittsburgh, Pittsburgh, Pennsylvania 15260
Juan José Mendoza Arenas, jum151@pitt.edu
University of Pittsburgh, Pittsburgh, Pennsylvania 15260
Nikita Gourianov, NIG134@pitt.edu
University of Pittsburgh, Pittsburgh, Pennsylvania 15260
Michael Lubasch, micahel.lubasch@cambridgequantum.com
Cambridge Quantum, 32 St. James Street, London, United Kingdom
Andrew J. Daley, andrew.daley@strath.ac.uk
University of Strathclyde, Glasgow, Scotland G4 0NG, United Kingdom
Dieter Jaksch (Presenter), dieter.jaksch@uni-hamburg.de
Universität Hamburg, Luruper Chaussee 149, Gebäude 69, D-22761 Hamburg, Germany and
Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

Topic: We propose developing quantum algorithms and quantum software for solving classical problems in computational science and engineering (CSE) with vast range of applications from computational fluid dynamics and weather forecasting to plasma physics, quantum chemistry, materials science and industrial design optimization. A large majority of these types of problems are formulated as either optimization problems and/or sets of partial differential equations (PDEs). Methods for translating these problems into a form that is suited to being solved efficiently on quantum computers are required. This will entail extending recently developed quantum algorithms to become applicable to CSE problems and should also lead to the invention of entirely novel quantum approaches that optimally utilize existing and future quantum hardware.

Challenges: Recently developed quantum algorithms promise efficient solutions to computational sub-problems that arise in the context of CSE. Sets of linear equations can be solved using the HHL approach [1], and hybrid quantum-classical variational algorithms for linear [2] and for nonlinear optimization problems [3] have also been developed. Exact diagonalization problems can be mapped onto Variational quantum eigensolvers [2]. Proof of principle demonstrations, for instance in simulating quantum fluids [3], were shown to only require a small number of qubits and have proven to be relatively noise tolerant.

Despite of this progress little is known on how these algorithms can be extended and adapted to classical problems in CSE and to understand where a quantum advantage might be expected. This work faces the following challenges:

(i) **Representability:** Solutions to classical problems like flow fields, pressure fields, temperatures, etc. need to be represented on quantum computers and must be encoded into quantum states efficiently. Finding efficient structures of quantum networks for achieving such encodings is in general an open question. To overcome this challenge, known properties of a solution can be exploited. For instance, it was recently shown that the nature of eddy interactions in turbulent fluid flows allows representing them efficiently with logarithmically small quantum registers and shallow quantum networks [4, 5].

(ii) **Classical optimization:** Hybrid optimization algorithms introduced above require classical optimization of quantum network parameters via feedback loops. The optimization landscape of highly expressive variational ansatzes contains large regions with vanishing gradients and local minima, which prevent efficient optimization. However, the quantum nature of the underlying optimization problem enables strategies to avoid such “barren plateaus” like the usage of classical shadows [6]. The structure of the optimization landscape of CSE problems is still largely unexplored and further research into improving classical optimization in hybrid algorithms is required.

(iii) **Problem specification:** CSE problems typically come with complex boundary conditions and design constraints that limit the range of possible solutions. In many cases it is known, in principle, how problem specifications can be encoded into a quantum algorithm. However, naïve approaches quickly lead to quantum networks that exceed the capabilities of current and near future quantum hardware. Thus, it will be important to find efficient encodings that optimally utilize the available quantum resources.

(iv) **Gaining a quantum advantage:** The exponential size of Hilbert space allows representing solutions to classical problems using only a logarithmically small number of qubits compared to storage in classical memory. In order to turn this favorable scaling into a quantum advantage, quantum algorithms need to scale polynomially in the number of qubits and overall better than classical algorithms. An automated a priori decision on the suitability of a given CSE problem to run on quantum hardware should be made with quantum hardware being utilized where an advantage can be expected.

Opportunities: Tackling the challenge of utilizing quantum algorithms in CSE problems requires different communities to come together. Importantly, progress in this area requires experts in CSE, who have a deep understanding of the structure of the problems and their solutions, and quantum algorithm experts to collaborate. Also, quantum hardware developers need to be involved to ensure that the quantum software is developed such that its demands can be met by near future hardware. See Ref. [7] for opportunities in aeroscience and engineering.

A particularly appealing example of where quantum hardware may be exploited successfully in CSE is the study of fluid flows, including compressible turbulent flow and combustion reactions. Quantum-inspired tensor network simulations for two- and three-dimensional systems show that quantum computing can reduce, by up to one order of magnitude, the necessary number of degrees of freedom to reproduce, with high accuracy, the exact turbulent dynamics obtained from state-of-the-art direct numerical simulation (DNS) [4]. This strategy has also been followed to study plasma physics [8] and establishes an appealing approach to solve problems of wide applicability such as the Boltzmann and Fokker-Planck equations [9]. Porting these approaches to quantum computers will at least provide a quadratic speed-up in grid size over standard classical simulations akin to the speed-up promised by the seminal Grover search algorithm.

Assessment: The ultimate goal in quantum algorithm development for CSE will be to gain a quantum advantage over classical algorithms in at least one of the possible figures of merit like, e.g., the number of required basic gate operations, the wall clock time, quantum memory requirements, or energy consumption. However, gaining such an advantage will crucially depend on progress in the development of quantum hardware (not part of this research).

Thus, success of research into quantum algorithms will be measured by assessing their scaling compared to well-established classical methods. For instance, classical simulations based on matrix product descriptions of specific CSE problems can provide bounds for the performance of corresponding quantum algorithms [5]. Initial studies indicate that gaining a quantum advantage should be possible for a quantum computer with around two- to three-hundred error-corrected qubits that can be operated upon fault tolerantly.

Importantly, variational quantum algorithms are often found to be capable of tolerating errors and results can be improved by error mitigation. This may lead to quantum advantages being obtained already on quantum devices that do not beat the error-correction threshold and operate fault tolerantly. Any such advantage would be highly dependent of the details of the quantum computing platform.

Timeliness or maturity: Quantum hardware developers promise quantum devices with several thousands of qubits by 2025 [10] and fault-tolerant devices by the end of the decade. Algorithms to be run on these devices must be developed contemporaneously to optimally utilize quantum hardware capabilities and realize the benefits early. Working on quantum codes now will also provide additional momentum to hardware developers and be instrumental in making hardware design decisions that ensure that a wide range of problems can be run beneficially on the hardware.

References:

[1] A.W. Harrow, A. Hassidim and S. Lloyd, *Quantum Algorithm for Linear Systems of Equations*, Phys. Rev. Lett. **103**, 150502 (2009).

[2] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P.J. Love, A. Aspuru-Guzik and J.L. O'Brien, *A variational eigenvalue solver on a photonic quantum processor*, Nat. Commun. **5**, 4213 (2014).

[3] M. Lubasch, J. Joo, P. Moinier, M. Kiffner and D. Jaksch, *Variational quantum algorithms for nonlinear problems*, Phys. Rev. A **101**, 010301(R) (2020).

[4] N. Gourianov, M. Lubasch, S. Dolgov, Q.Y. van den Berg, H. Babaee, P. Givi, M. Kiffner and D. Jaksch, *A quantum-inspired approach to exploit turbulence structures*, Nat. Comput. Sci. **2**, 30 (2022).

[5] D. Jaksch, P. Givi, A.J. Daley and T. Rung, *Variational quantum algorithms for CFD*, AIAA J. **61**, 1885 (2023).

[6] S.H. Sack, R.A. Medina, A.A. Michailidis, R. Kueng and M. Serbyn, *Avoiding Barren Plateaus Using Classical Shadows*, Phys. Rev. X Quantum **3**, 020365 (2022).

[7] P. Givi, A.J. Daley, D. Mavriplis, and M. Malik, *Quantum Speedup for Aeroscience and Engineering*," *AIAA J.*, **58**, 3715 (2020)

[8] E. Ye and N.F.G. Loureiro, *Quantum-inspired method for solving the Vlasov-Poisson equations*, Phys. Rev. E **106**, 035208 (2022).

[9] A.M.P. Boelens, D. Venturi and D.M. Tartakovsky, *Parallel tensor methods for high-dimensional linear PDEs*, J. Comput. Phys. **375**, 519 (2018).

[10] See e.g. IBM Unveils 400 Qubit-Plus Quantum Processor and Next-Generation IBM Quantum System Two (accessed 27.4. 2023)

The Next Mile in Quantum Software

Pranav Gokhale[1, *]
[1] *Super.tech, a division of Infleqtion*

TOPIC

Compilation; Codesign and Integration

I. CHALLENGE

Standards such as OpenQASM 3[1], QIR [2], and OpenPulse [3]—as well as their downstream open-source software packages—have made important progress toward closing the gap between quantum software and hardware. For example, OpenQASM 3 has rich syntax and primitives for expressing device-level timing constraints, pulse calibrations, and classical control flow. This is a significant advance from the earliest quantum assembly languages [4] which did not have the benefit of multiple readily-accessible quantum hardware platforms.

Despite this progress, our experience is that few quantum hardware platforms leverage these standards or their downstream open-source software packages beyond basic usage and interfacing. Why? We identify three exemplar challenges that force quantum hardware developers to develop ad hoc quantum software rather than integrating more deeply with broader standards and quantum software libaries. These challenges should be addressed on the next (though certainly not *last*) mile in quantum software.

1. **Virtual phase tracking**. On many systems, especially superconducting and trapped ions, $R_z(\theta)$ gates can be realized *virtually* at zero cost by deferring implementation to subsequent physical gates [5]. Existing quantum software largely lacks support for tracking virtual phase accumulation, instead deferring this aspect to proprietary and vendor-specific control software. This presents missed opportunities for deeper compiler integration. Consider: virtual phase don't commute with entangling gates such as the B (Berkeley) gate or Mølmer-Sørensen XX-interaction. For these gates, a Phased MicroWave (PMW) decomposition [6] is needed whereby virtual phases are zeroed prior to the entangler. Mainstream quantum software must account for virtual phase tracking to enable exploration of emerging gate protocols. In addition, in cases like iSWAP, virtual phase on one qubit commutes to the *other* qubit, as depicted below. This again needs software support. In our experience, gate protocol development has been stymied by the lack of virtual phase tracking in leading quantum software.

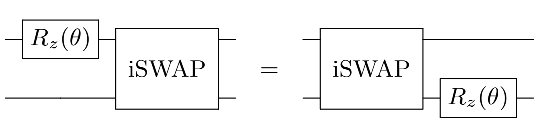

2. **Complex Scheduling**. As quantum software has advanced closer to hardware, a new set of constraints have emerged owing to the device physics of gate scheduling. We provide two examples. (a) Researchers at the Berkeley Advanced Quantum Testbed have found remarkable success in calibrating special composite gates like $\{CZ(q_0, q_1), CZ(q_2, q_3)\}$ which exhibits lower crosstalk/spectator error relative to independent pulses for the two constituent gates [7]. (b) On most neutral atom hardware, single-qubit R_{xy} gates must be executed *globally* on every qubit at the same time: no single-site R_{xy} is possible [8]. Both of these scheduling challenges are currently being addressed with ad hoc post-compilation scripts developed by hardware labs. However, these low-level scheduling requirements should be moved into the compiler where they can achieve greater optimality by integration with other compiler passes.

3. **Software maintenance demands**. Quantum software package owners are often in the unenviable position of "dependency hell." On one hand, it is extremely useful and productive to import leading open-source quantum software packages to avoid re-inventing the wheel. However, these leading packages often have conflicting dependencies, frequent backwards-incompatible upgrades, and complicated installation schemes. As a result, we find that many leading software packages developed by quantum researchers become unusable within $\sim$ two years of initial release. This recurring obsoletion process both hampers scientific progress and reduces trust in the reliability of quantum software.

II. OPPORTUNITY

Challenges #1 and #2 are specific instances of a broader set of emerging challenges that are coming into view as quantum software becomes more and more aware of idiosyncracies in underlying quantum hardware. There is an opportunity at hand to ensure that development of low-level compilation does not occur in separate vacuums for each hardware platform, but rather through shared infrastructure that moves the entire field forward—including for yet-to-be-discovered qubit types. For example, virtual phase tracking software for PMW could enable both operation of a B gate on superconducting systems as well as the operation of exotic two-qubit interactions on neutral atom systems. The first step to addressing these challenges would be to convene stakeholders from all leading qubit types to understand what aspects of their quantum software stack are not currently covered by leading tools and standards.

Challenge #3 (Software maintenance demands) is not

* pranav.gokhale@infleqtion.com

a scientific challenge but rather a mechanical one that hampers scientific progress. We believe there is a timely opportunity for industry engagement to play a role in maintaining, upgrading, and enhancing software that is first conceived by researchers at national laboratories. We believe this division of labor is also better compatible with incentives and expertise: often researchers at national laboratories are experts at 0-to-1 inventions whilst industry incentives are better aligned for 1-to-N scaling. For example, industry can develop software tools similar to the Github Dependabot (which automatically upgrades software package dependencies based on security updates), specialized to quantum software dependencies. The first step to addressing these challenges would be to identify 3–5 software packages developed by national labs that are appropriate for "graduating" to maintenance either by an industry player or by an open-source consortium. An exemplar of the latter is the XACC framework [9], which is now part of the Eclipse organization.

III. ASSESSMENT

For the first two Challenges, the primary metric of success would be deeper adoption of quantum standards and open-source libraries by leading hardware developers. As a corollary, hardware developers, especially from research institutions should be naturally incentivized to contribute updates back upstream. This should be measured by cross-pollination in quantum software across qubit technologies: for example, how often are low-level compiler developments from superconducting laboratories used by trapped ion groups—and vice versa? Potential solutions should be evaluated by adoption.

For the software maintenance Challenge, a simple metric would be: what fraction of open-source software packages developed by researchers are usable (i.e. can `pip install` and successfully run all example scripts and unit tests) two years after their initial release? As a corollary, we should see greater community confidence in open-source quantum software, which could be measured through mechanisms such as Unitary Fund's annual Quantum Open Source Software Survey [10].

IV. TIMELINESS OR MATURITY

We first emphasize that there have been significant improvements to quantum software over the past decade. Every year, tens of thousands of students and researchers submit programs to quantum computers using Qiskit, Q#, Braket, and Cirq—an enormous step forward from just ten years ago.

The opportunity to carry forward this momentum—with even deeper integration with quantum hardware—builds on top of the existing success of quantum software. As referenced earlier, quantum software integration of features like virtual phase tracking will enable operation of novel two-qubit gates with higher fidelity, which remains one of the key challenges on the path to utility-scale quantum computation. In addition, the opportunity to share optimizations across qubit modalities would be a high-ROI venture that can be uniquely enabled by advances in quantum software.

Finally, the identified challenges pertaining to software maintenance are particularly timely, since the research community has increasingly embraced artifact-evaluation and code-sharing platforms like GitHub and Zenodo. There is a window of opportunity to ensure that this uptick in researcher-developed software packages becomes contributes multiplicatively to subsequent research and development.

[1] A. Cross, A. Javadi-Abhari, T. Alexander, N. De Beaudrap, L. S. Bishop, S. Heidel, C. A. Ryan, P. Sivarajah, J. Smolin, J. M. Gambetta, *et al.*, Openqasm 3: A broader and deeper quantum assembly language, ACM Transactions on Quantum Computing **3**, 1 (2022).

[2] QIR Alliance, *QIR Specification* (2021), also see `https://qir-alliance.org`.

[3] D. C. McKay, T. Alexander, L. Bello, M. J. Biercuk, L. Bishop, J. Chen, J. M. Chow, A. D. Córcoles, D. Egger, S. Filipp, *et al.*, Qiskit backend specifications for openqasm and openpulse experiments, arXiv preprint arXiv:1809.03452 (2018).

[4] K. M. Svore, A. V. Aho, A. W. Cross, I. Chuang, and I. L. Markov, A layered software architecture for quantum computing design tools, Computer **39**, 74 (2006).

[5] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient z gates for quantum computing, Physical Review A **96**, 022330 (2017).

[6] J. Chen, D. Ding, C. Huang, and Q. Ye, Compiling arbitrary single-qubit gates via the phase-shifts of microwave pulses, arXiv preprint arXiv:2105.02398 (2021).

[7] A. Hashim, R. Rines, V. Omole, R. K. Naik, J. M. Kreikebaum, D. I. Santiago, F. T. Chong, I. Siddiqi, and P. Gokhale, Optimized swap networks with equivalent circuit averaging for qaoa, Physical Review Research **4**, 033028 (2022).

[8] T. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, *et al.*, Multi-qubit entanglement and algorithms on a neutral-atom quantum computer, Nature **604**, 457 (2022).

[9] A. J. McCaskey, D. I. Lyakh, E. F. Dumitrescu, S. S. Powers, and T. S. Humble, XACC: a system-level software infrastructure for heterogeneous quantum–classical computing, Quantum Science and Technology **5**, 024002 (2020).

[10] Unitary Fund, *Quantum Open Source Survey 2022* (2022).

Speeding up quantum computers with mid-circuit measurements

Alexey Gorshkov, University of Maryland, Email: gorshkov@umd.edu
Andrew Childs, University of Maryland, Email: amchilds@umd.edu
Nathan Schine, University of Maryland, Email: nschine@umd.edu

Topic: This position paper addresses the first research area: *"The quantum software stack and fundamental quantum computer science and algorithms research"*. In particular, this position paper proposes to study how fast mid-circuit measurements and fast classical feedback can speed up quantum computers. In particular, mid-circuit measurements will allow the *compilation* of quantum algorithms into significantly-lower-depth circuits than is possible with purely unitary constructions.

Challenge: Qubit connectivity plays an important role in how quickly one can run a given quantum algorithm on a quantum computer. While any algorithm would run the fastest in an architecture that features all-to-all connectivity, such connectivity is technically impossible to achieve for large qubit numbers. Therefore, a typical quantum computing architecture, such as for example superconducting hardware, features an array of qubits with interactions obeying some sort of locality constraints, such as, for example, the ability to implement gates only between neighboring qubits. The unitary implementation of quantum algorithms in such spatially local architectures is constrained by Lieb-Robinson bounds (see e.g. [1]). For example, if one needs to implement a gate between two qubits a distance r apart, such a gate would necessarily take time t linear in r as one needs to bridge the distance between the two qubits. The resulting constraints result in significant slow-downs (see e.g. [2]) in the implementation of quantum algorithms compared to what one could do with—unfortunately unachievable—all-to-all connectivity. Given the limited coherence of available qubits, such slow-downs pose an important challenge for the goal of scaling quantum systems up and implementing quantum algorithms with high fidelity.

Opportunity: While unitary dynamics is constrained by Lieb-Robinson bounds, if one is able to do fast mid-circuit measurements and apply fast classical feedback—as has been already demonstrated for trapped-ion [3–6], superconducting [7], and neutral-atom [8] architectures—then Lieb-Robinson bounds essentially do not apply, provided one can make sufficiently many measurements [9]. In particular, a quantum repeater protocol allows one to send quantum information across distance r—and hence apply a quantum gate between two qubits separated by distance r—in time that does not depend on r [10]. Similarly, one can apply an unbounded quantum fanout gate, which is an important ingredient in Shor's algorithm, in constant time [10]. We have also shown that fast mid-circuit measurements allow one to route quantum information faster across a quantum computing architecture [11]. We [12] and others (see e.g. [13–15]) have also shown that fast mid-circuit measurements allow one to quickly prepare a variety of many-body entangled states. The above-mentioned results notwithstanding, it is still very much an open question to what degree mid-circuit measurements allow one to speed up the implementation of any given quantum algorithm (i.e. any given unitary). In fact, even for the case of relatively simple unitaries that involve only permutations of qubits (as in qubit routing), there are huge gaps between the available mid-circuit-measurement-based algorithms and the corresponding lower bounds on the implementation time [11]. In this position paper, we propose to study how quickly one can implement a given quantum algorithm (i.e. a given unitary) on a given architecture with the help of mid-circuit measurements. The goal is to derive both fast protocols and matching lower bounds.

Some experimental architectures may involve qubits coupled to optical or microwave cavities that are used for fast readout. In fact, one of us (Schine) is building a neutral-atom tweezer-array experiment, in which fast mid-circuit measurements will be carried out with the help of a cavity. In fact, one could argue that using a cavity for readout is necessary if one wants truly fast mid-circuit measurements. Since it is very difficult to design an experiment where each qubit is coupled to its own cavity (or cavity mode), it is reasonable to expect that many qubits will be coupled to the same cavity mode, as in the experiment that PI Schine is building. In that case, one is not able to simultaneously measure every qubit (via the cavity), which places a limitation on the types of mid-circuit measurements one can perform per unit time. At the same time, in such cavity-based architectures, one is able to make collective (global) measurements on any desired subset of qubits coupled to the same cavity mode, which endows such architectures with additional capabilities. Therefore, in addition to studying the more commonly assumed scenario where any subset of qubits can be measured independently and simultaneously, we will also extend our study of mid-circuit measurements to the situation where qubits are broken into groups (each group is coupled to its own cavity mode), and one is allowed to make in unit time a collective measurement on any desired subset of any given group, and this can be done in parallel on every group.

While we will start our analysis without taking imperfections into account, we will eventually incorporate imperfections and error correction.

Assessment: The project would be successful if, for any given architecture, we develop fast mid-circuit-measurement-based protocols and matching lower bounds for a wide range of most commonly used unitaries and ideally for a range of entire quantum algorithms. The project would be wildly successful if, for any given unitary, we can provide mid-circuit-measurement-based protocols and lower bounds on the implementation time that are not too far from each other.

Timeliness: Given the recent experimental demonstration of fast mid-circuit measurements and fast classical feedback in trapped-ion [3–6], superconducting [7], and neutral-atom [8] architectures, now is the perfect time to figure out in detail how much mid-circuit measurements can speed up the implementation of quantum algorithms. Indeed, as mentioned above, one of us (Schine) is building an experimental system that features a particularly interesting mid-circuit-measurement capabilities that call for being studied theoretically.

[1] M. C. Tran, C.-F. Chen, A. Ehrenberg, A. Y. Guo, A. Deshpande, Y. Hong, Z.-X. Gong, A. V. Gorshkov, and A. Lucas, Phys. Rev. X **10**, 031009 (2020).
[2] A. Bapat, A. M. Childs, A. V. Gorshkov, and E. Schoute, PRX Quantum **4**, 010313 (2023).
[3] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, et al., arXiv:2003.01293 (2020).
[4] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, et al., arXiv:2107.07505 (2021).
[5] D. Zhu, G. D. Kahanamoku-Meyer, L. Lewis, C. Noel, O. Katz, B. Harraz, Q. Wang, A. Risinger, L. Feng, D. Biswas, et al., arXiv:2112.05156 (2021).
[6] M. Foss-Feig, A. Tikku, T.-C. Lu, K. Mayer, M. Iqbal, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, A. Hankin, et al., arXiv:2302.03029 (2023).
[7] A. D. Córcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Minev, J. M. Chow, and J. M. Gambetta, Phys. Rev. Lett. **127**, 100501 (2021).
[8] T. M. Graham, L. Phuttitarn, R. Chinnarasu, Y. Song, C. Poole, K. Jooya, J. Scott, A. Scott, P. Eichler, and M. Saffman, arXiv:2303.10051 (2023).
[9] A. J. Friedman, C. Yin, Y. Hong, and A. Lucas, arXiv:2206.09929 (2022).
[10] P. Pham and K. M. Svore, Quantum Info. Comput. **13**, 937 (2013).
[11] D. Devulapalli, E. Schoute, A. Bapat, A. M. Childs, and A. V. Gorshkov, arXiv:2204.04185 (2022).
[12] Z. Eldredge, L. Zhou, A. Bapat, J. R. Garrison, A. Deshpande, F. T. Chong, and A. V. Gorshkov, Phys. Rev. Research **2**, 033316 (2020).
[13] N. Tantivasadakarn, R. Thorngren, A. Vishwanath, and R. Verresen, arXiv:2112.01519 (2021).
[14] R. Verresen, N. Tantivasadakarn, and A. Vishwanath, arXiv:2112.03061 (2021).
[15] N. Tantivasadakarn, A. Vishwanath, and R. Verresen, arXiv:2209.06202 (2022).

New Methods for Assessing Progress in the Development of Quantum Systems

Luke C. G. Govia,[1,*] David McKay,[2] Seth Merkel,[2] Andrew Wack,[2] and Travis L. Scholten[2]

[1]*IBM Quantum, IBM Almaden Research Center, San Jose, California 95120, USA*

[2]*IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA*

Topics: Applications, algorithms, compilation, error correction/mitigation, codesign, integration

I. Challenges

We identify 4 primary challenges which must be overcome to support the further maturation of quantum computing platforms.

A. Creating new metrics for the performance of integrated quantum/classical compute systems

A quantum-centric supercomputer (QCSC) [1] promises a tighter integration of quantum and classical compute resources. New methods to characterize performance and data flow between disparate compute resources need to be developed which by necessity will benchmark the full compute stack, or a large sub-span of it. First steps in this direction such as quantum volume (and variants) [2, 3] and CLOPS (Circuit Layer Operations Per Second) [4] study lower layers of the stack: e.g. hardware, compilation, system architecture, and control electronics. Future metric design must incorporate the idea of a QPU as a co-compute unit.

B. Formalizing benchmark development

There is no one-size-fits-all benchmark for quantum systems, and benchmarks that capture performance across the axes of scale, quality, and speed are necessary. Further, how a given benchmark is actually run in practice, e.g. the gateset used in randomized benchmarking [5], or the statistical tests in quantum volume [2, 3], can make a dramatic difference in the results. For this reason, a benchmarking suite must come equipped with a standard set of operating procedures and rules [6]. Finally, the community needs a more integrated perspective on how to interpret and apply the results of a given benchmark, both within and outside its defined context. For example, how well can benchmarks capture performance outside of their specific domain, or can sub-system benchmarks predict holistic performance?

C. Developing application-level performance benchmarks

Application-level benchmarks must enable end-users to make meaningful inferences about how well a given system would do in tackling their particular problem, *without* relying solely on running test instances of that problem. Moreover, the gap in prior work [7–10] is a lack of measures applicable to applications *qua* applications. If the algorithms and application development community can find commonalities between disparate application domains, e.g. in circuit sub-routines, then cross-application benchmarking may be possible. Noise tailoring and codesign are useful, but cannot be applied to each application in a fully application-specific way, so identifying the trade-offs here is also crucial.

The prediction of application performance relies on predictive/integrated noise characterizations of subsystem characterizations that can be combined into an error map for an entire device or QC stack. Further, there is a need to understand the interplay between "classical errors" in control, compilation, etc with "quantum errors" at the device level, and identify when the total error is more than the sum of its parts. Finally, the community currently lacks efficient and accurate ways of validating characterized models, without which performance predictions cannot be trusted to be accurate.

D. Analyzing the impact and limits of error mitigation

The rapid recent development in error mitigation (EM) techniques [11, 12] fundamentally changes the impact of noise on algorithms or applications. EM aims to reduce the error in a circuit at the cost of a computational overhead coming from the need to run an ensemble of circuits. As EM will be one of the workhorse techniques the community uses for the foreseeable future, understanding how to incorporate it into assessing the performance of quantum systems is necessary. Further, EM techniques can require detailed noise characterization, motivating the study into accurate, efficient, and most importantly scalable characterization protocols. Despite its maturing development, there are many open questions in the interplay of noise and EM, such as: what noise processes cannot be corrected this way, and what families of circuits are well-suited to stress-testing the effectiveness of EM techniques, i.e. form benchmarks of EM itself?

II. Opportunity

The identified challenges exist at the intersection of all stakeholders in the QC stack, from end-users to hardware developers, and necessarily includes benchmark development and standardization groups. One of the best ways to address the identified challenges is to facilitate communication and collaboration between these groups, so that priorities at all levels of the stack can be identified and incorporated into benchmark development.

There is a need for more collaboration between the Computer Science & High-Performance Computing community and the quantum community, so the latter can leverage strengths of the former when it comes to creating new holistic metrics for quantum-classical systems at scale (Challenge A). Further, there is a need for more

* Corresponding author: lcggovia@ibm.com

collaboration between researchers and standards-setting bodies, so that the deliberations of the latter are leveraging the expertise of the former (Challenge B).

Several new techniques and tools need to be developed, including those to study the applicability of application-based benchmarks in other contexts, to understand how errors combine across the QC stack, to enable validation of noise models (Challenge C), and to assess the limits of error mitigation techniques (Challenge D). The importance of techniques to learn *detailed* descriptions of the error mechanisms at all levels of the stack cannot be underestimated. Such detailed information directly impacts not only hardware development, but also software design choices, compilation and error mitigation strategies, and even algorithm co-design.

III. Assessment

The impact of benchmarking and characterization is more clearly demonstrated in their absence, where development moves forward without a clear goal or direction in mind, and progress ends up being more accidental than intentional. The classic example in QC is the development of randomized benchmarking, which enabled estimation of the impact of a wide variety of error sources on quantum gates, and has guided the rapid improvement in gate fidelity across all platforms over the last decade.

Success of a benchmarking or characterization protocol should be directly tied to how it guides development towards improved performance, either of a specific component, holistically of a larger subset of the stack, or of an application domain. While they may also be useful to compare processors or implementations, the field is not in a position where standardization is mature enough that cross-platform comparison is a sufficient use case for benchmarks. To take a specific example, quantum volume has been widely adopted as a comparative benchmark, even across physical platforms. However, one could argue the major success of quantum volume has been that it highlighted the importance of aspects of the stack beyond quantum processor performance, e.g. compilation and parallelization, which has led to rapid development on these fronts.

The natural conclusion of this line of thought is that the more detailed the output of a characterization protocol, the more successful it will be. To some extent, this is true, but benchmarks must also be judged on how easily they can be implemented, how easily their output can be interpreted, and on their scalability. For example, it is along these axes that randomized benchmarking outperforms the more detailed process tomography. Ultimately, a successful characterization or benchmarking protocol is one that enables the user to drive meaningful progress as their systems grow in size, and which is broadly adopted by the community.

One of the key resources needed to make progress on these challenges is access to quantum systems with advanced hardware and software capabilities. Without such access, addressing the challenges may become very theoretical in nature, and will not materially help the community create a path forward for assessing progress in the development of quantum systems.

IV. Timeline and maturity

One of the primary reasons to address the challenges above is that the rapid expansion of the quantum community has led to a proliferation of benchmarks and metrics, which has enabled tremendous learning, but created confusion about nomenclature and what needs to be measured and assessed to evaluate progress. Governments and other stakeholders have keen interest to get a handle on the performance of quantum computers as new national initiatives are launched and the competitive landscape evaluated. In addition, the scale and complexity of cloud-accessible quantum systems has grown, from 5 qubits in 2016 to 433 in 2023. The availability of advanced systems and software capabilities enables studying these challenges on deployed systems, and also has rendered their solution more timely and urgent.

[1] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, Journal of Applied Physics **132**, 160902 (2022).

[2] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, Phys. Rev. A **100**, 032328 (2019).

[3] C. H. Baldwin, K. Mayer, N. C. Brown, C. Ryan-Anderson, and D. Hayes, Quantum **6**, 707 (2022).

[4] A. Wack, H. Paik, A. Javadi-Abhari, P. Jurcevic, I. Faro, J. M. Gambetta, and B. R. Johnson, arXiv:2110.14108 (2021).

[5] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Phys. Rev. A **77**, 012307 (2008).

[6] M. Amico, H. Zhang, P. Jurcevic, L. S. Bishop, P. Nation, A. Wack, and D. C. McKay, arXiv:2303.02108 (2023).

[7] D. Mills, S. Sivarajah, T. L. Scholten, and R. Duncan, Quantum **5**, 415 (2021).

[8] T. Lubinski, S. Johri, P. Varosy, J. Coleman, L. Zhao, J. Necaise, C. H. Baldwin, K. Mayer, and T. Proctor, IEEE Transactions on Quantum Engineering (2023).

[9] R. Blume-Kohout and K. C. Young, Quantum **4**, 362 (2020).

[10] S. Martiel, T. Ayral, and C. Allouche, IEEE Transactions on Quantum Engineering **2**, 1 (2021).

[11] E. v. d. Berg, Z. K. Minev, A. Kandala, and K. Temme, arXiv:2201.09866 (2022).

[12] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2020) pp. 306–316.

Quantum Algorithms for Efficient Time-Dependent Quantum Dynamics

PI: Niranjan (Niri) Govind*, Physicist, Pacific Northwest National Laboratory
Co-PI: Bo Peng, Computational Scientist, Pacific Northwest National Laboratory
Co-PI: Stephen Gray, Senior Scientist, Argonne National Laboratory
Co-PI: Yuri Alexeev, Principal Project Specialist, Argonne National Laboratory
niri.govind@pnnl.gov

Call: Position Paper for Workshop on Quantum Computing and Networking
Topics: Algorithms and Compilation

Challenge: Time-dependent quantum dynamics is essential for a fundamental understanding of the equilibrium and out-of-equilibrium dynamics of many-body quantum systems. However, due to the curse of dimension, quantum many-body problems require a Hilbert space that grows exponentially with the number of degrees of freedom, making these problems intractable with classical approaches. By leveraging the power of superposition and entanglement to efficiently sample the exponentially large Hilbert space of these problems, quantum computers offer a solution out of this conundrum. As suggested by Benioff[1] and Feynman,[2] it is only natural to simulate quantum systems using quantum computers and, going further, one may argue that quantum time dynamics is an ideal candidate in the continuing pursuit for quantum supremacy. The first step towards this quest is to formulate the quantum many-body problem as a Hamiltonian, which can then be exponentiated (and Trotterized, if needed) to perform the time evolution of interest. The time propagator can always be encoded in a qubit representation, and the resulting quantum circuits can be run on a quantum computer. However, these steps do not guarantee the achievement of quantum advantage over classical computing. It is well known that quantum circuits representing quantum time dynamics grow with increasing time simulations. The circuits are either too large and/or too deep to execute even on a fault-tolerant quantum device, given time limitations. Nevertheless, it is possible to effectively compress certain types of quantum circuits representing integrable models.[3,4] Generalizing these compression techniques to non-integrable models is an ongoing challenge, which will allow one to study the quantum time dynamics of broader classes of quantum many-body systems and potentially achieve quantum advantage in the future. More broadly, quantum time dynamics in combination with error mitigation and correction techniques could help catalyze a continuous transition towards the fault-tolerant regime.

Opportunity: We have recently demonstrated[4] a promising quantum circuit compression technique using the foundational Yang-Baxter equation (YBE). The YBE[5,6] is a consistency relation that allows one to factorize three-body interactions into a sequence of pairwise interactions under specific conditions. It has been exploited to solve certain classes of quantum many-body problems.[7,8] By taking advantage of the YBE, we were able to compress and produce a shallow quantum circuit for efficient quantum time dynamics simulations of special cases of the Heisenberg spin model on real quantum computers.[4] With this approach, the resulting constant depth circuit is independent of time and step size. The compressed circuit a linear function of the system size, and the number of CNOT gates only scales quadratically with system size, allowing for long time dynamics simulations. To demonstrate the efficacy of the method, we performed quantum time dynamics simulations of three and five spins on an IBM quantum computer and compared the dynamics from both compressed and uncompressed quantum circuits. For the first time, our results confirmed the superiority of the YBE formulation in performing dynamics for many steps and connected our work to the broader and deeper context of the YBE duality and integrable quantum computation. We have also developed an open-source algebraic compiler (QuYBE) to compress quantum circuits. The compiler[9] is a general YBE-based quantum circuit compression algorithm that can perform compression for arbitrary N-qubits. QuYBE is the first step towards making this approach to the broader community. The QuYBE compiler is available at https://github.com/ZichangHe/QuYBE.

Going beyond the YBE approach, other algebraic expressions, such as the Cartan decomposition of Lie algebra generated by the Hamiltonian[10,11], can be explored. These circuit ansatzes can be used to replace the original circuit fragment employed in the quantum time dynamics circuit with sufficient accuracy to facilitate circuit compression, compilation, and optimization, achieving sufficiently high fidelity of the generated state. One may also construct suitable approximations to study the quantum time dynamics of fully non-integrable models and ones that can be adiabatically (or perturbatively) extended from integrable models. Other opportunities include quantum imaginary time evolution of open quantum systems, the effect of noise on the quantum time dynamics of integral models, and quantum time dynamics within partitioned (or embedded) subspaces of a quantum many-body system. Additionally, as an aside, the quantum time dynamics of integrable models are good candidates since they can be simulated in a scalable way on classical computers, and when combined with YBE-based circuit compression, can provide a large tuning space for evaluating the efficacy of quantum hardware platforms.

Assessment: Our YBE technique, along with its generalization towards a robust quantum compiler, represents a comprehensive set of approaches that can be applied to any partially or fully compliant quantum circuit. Combining this approach with a recently developed Krylov-subspace-based approach[12] for time propagation could yield even more efficiencies. Our success metric will therefore include the following elements:

- Comprehensive evaluation of the proposed YBE technique, including a comparison with existing quantum dynamics simulation methods, an analysis of the scalability and error mitigation capabilities, and a discussion on the integration of classical methods for hybrid quantum-classical approaches. This should be done alongside the assessment of common long-term decomposition and short-term variational quantum algorithms for quantum dynamics[13].
- Demonstrations of the effectiveness of the YBE technique in various applications, such as chemistry, physics, materials science, drug discovery, and optimization problems, by showcasing efficient quantum simulations of electron dynamics in molecules with the order of 50 to 100 electrons, or the non-equilibrium dynamics of lattice spin Hamiltonian with $\mathcal{O}(10^2)$ sites.
- Exploration of the combination of error mitigation techniques and quantum control in the proposed YBE technique, demonstrating the progression towards the fault-tolerant regime and highlighting the potential applicability of the method to practical, real-world problems.

Timeliness or maturity: Despite early suggestions that quantum time dynamics would be a key application for quantum computing, actual studies beyond model systems have been lacking. However, recent advancements in quantum hardware, error mitigation techniques, and noise reduction methods have made it possible to perform more complex quantum simulations and tackle larger systems. As a result, our proposed quantum algorithms and quantum compilers for efficient quantum time dynamics can capitalize on these timely improvements. Furthermore, our proposed approaches, combined with improved software tools and optimization techniques, puts us in an excellent position to study interdisciplinary quantum dynamics problems in the NISQ era. Therefore, the potential impact of successfully simulating quantum time dynamics is substantial. It could enable the design of new materials, understanding new physical phenomena in the condensed phase, reaction mechanisms, the discovery of new drugs, and provide a deeper understanding of fundamental physical and chemical processes. As the NISQ-era quantum computers and algorithms continue to improve, including our proposed algorithmic approaches for quantum time dynamics, we anticipate significant progress and a reshaping of this area and its applications.

References: [1] P. Benioff. *J. Stat. Phys.* 22, 5 1980. [2] R. P. Feynman. *Int. J. Theor. Phys* 21, 6/7 1982. [3] L. Bassman Oftelie et al. *Mater. theory* 6, 1 2022. [4] B. Peng et al. *Phys. Rev. A* 106, 1 2022. [5] C. N. Yang. *Phys. Rev. Lett.* 19, 23 1967. [6] R. J. Baxter. *Ann. Phys.* 70, 1 1972. [7] M. Jimbo. *Int. J. Mod. Phys. A* 4, 15 1989. [8] B. Sutherland. *J. Math. Phys.* 12, 2 1971. [9] S. Gulania et al. *2022 IEEE/ACM 7th Symposium on Edge Computing (SEC)*. IEEE. 2022. [10] E. Kökcü et al. *Phys. Rev. A* 105, 3 2022. [11] E. Kökcü et al. *Phys. Rev. Lett.* 129, 7 2022. [12] C. L. Cortes et al. *Phys. Rev. A* 106, 4 2022. [13] A. Miessen et al. *Nat. Comput. Sci.* 3, 2023.

Reconciling the tension between generality and specificity in quantum computing benchmarks with a hierarchical, application-oriented approach

Peter Graf (peter.graf@nrel.gov)
Computational Science Center, National Renewable Energy Laboratory

Topic: *Applications and models*

Challenge: *Fair and unifying yardsticks for QC*.
Application-oriented benchmarks have long had an important role to play in developing and scaling computational methods (e.g., LINPACK, MNIST). Given the state of affairs in quantum computing--namely, competing computational models and disparate, "work-in-progress" hardware--the need is arguably even more compelling here. And indeed, many benchmarks have been proposed and implemented (see [1-3], and many more).

There is an essential tension, however, specific to quantum computing benchmarks, between generality and specificity, and it relates to the ongoing problem of reconciling the different models of quantum computation and the many different and rapidly developing hardware instantiations of those models. On one hand, a good benchmark allows for a true test of a particular combination of machine and software. In quantum computing, this requires, to date, careful considerations stretching from mathematical formulation to the details of the hardware. On the other hand, we would like to construct "general" benchmarks that measure general capabilities. The problem, currently, is that the divergent combinations of models and hardware make the construction of general benchmarks fundamentally "unfair", as they will end up either being too general to be relevant, or too specific to be truly "model/hardware agnostic".

Opportunity: *Unity through a hierarchy of application-inspired benchmarks*.
Benchmarks come in many shapes and sizes. Some are calibration oriented, some represent "kernels'' from which algorithms can be built, and some involve "canonical" applications of VQE (to, e.g., diagonal Hamiltonians) or QAOA (to, e.g., MaxCut). Few go beyond this to true "exemplar" applications (e.g., using MaxCut to maximize traffic flow in real-time). The field has yet to coalesce around a structure that captures the full hierarchy of benchmarks and how they might be profitably used together to better target both research and attempts at quantum advantage.

We hypothesize there is an opportunity to "organize'' QC benchmarks in a hierarchical way that both entails a structure to the "space of benchmarks" and helps resolve the tension between computational-model and/or hardware agnosticism and the reality of current quantum computers.

Assessment: *Hypothetical high-level application benchmarks/exemplars*.
[1-3] and many others flesh out a range of "lower-level" benchmarks, from calibration to VQE and QAOA on model problems. There seems to be a lack of research pushing from the other end, i.e., from the starting point of *real* applications domains, for example, materials chemistry, power systems planning and control.

Often there is complexity theoretic work that says when a problem is hard. The value of application benchmarks, even within a certain complexity class, is that addressing a concrete problem clarifies the difficulty greatly [4] and is thus very helpful for figuring out when QC might be the best choice.

Here are three specific examples of "high-level" benchmarks:

- The Anderson Impurity Model, a system of highly correlated electrons postulated to elucidate many of the properties leading to quantum advantage for materials simulation. This model deliberately targets the “hardest” problem you can try to solve.
- Stochastic power systems planning and optimization. Many papers now study QC applied to *deterministic* power systems problems, and many claims are being made about the applicability of QC to this area. Two points: 1) there is no accepted set of benchmark QC formulations in this specific area; it would be extremely useful if there were; 2) there does not seem to be much QC research on *stochastic* programs, which are increasingly important as we transition to clean energy systems, where the primary resources (wind and solar) are fundamentally stochastic.
- Bond dissociation of TiH [4]. This paper makes the point that assessment of quantum readiness was only revealed via a detailed "implementation" of VQE (all but running it on real hardware). Only then could traditional benchmarks be utilized to complete the assessment.

Benchmarks that start with real target problems are important for assessing capabilities and requirements across model and machine classes.

Timeliness or maturity: *NISQ is ready for "full-stack" benchmarks.*
Benchmarks relying on fault-tolerant QC are not that useful yet. But the NISQ benchmark landscape is murky, because NISQ machines are murky. However, as we get closer to NISQ machines of practical size, it would accelerate progress to be able to assess, through benchmarks, their potential. I am suggesting an approach to benchmarks that 1) explicitly addresses the fact that there are not only competing machines but competing underlying models and 2) openly invites domain specialists to participate, because, given the zoo of models and machines, success on real applications is what we ultimately care about (and are finally able to consider).

References

[1] Cross, Andrew W., Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta. "Validating quantum computers using randomized model circuits." Physical Review A 100, no. 3 (2019): 032328.
[2] Lubinski, Thomas, Sonika Johri, Paul Varosy, Jeremiah Coleman, Luning Zhao, Jason Necaise, Charles H. Baldwin, Karl Mayer, and Timothy Proctor. "Application-oriented performance benchmarks for quantum computing." IEEE Transactions on Quantum Engineering (2023).
[3] Amico, Mirko, Helena Zhang, Petar Jurcevic, Lev S. Bishop, Paul Nation, Andrew Wack, and David C. McKay. "Defining standard strategies for quantum benchmarks." arXiv preprint arXiv:2303.02108 (2023).
[4] Clary, Jacob M., Eric B. Jones, Derek Vigil-Fowler, Christopher Chang, and Peter Graf. "Exploring the scaling limitations of the variational quantum eigensolver with the bond dissociation of hydride diatomic molecules." Int. J. Quantum Chem. 2023, e27097. https://doi.org/10.1002/qua.27097.

Noise-Aware Visual Recommendation System for Quantum Application Developers

Qiang Guan, Kent State University, qguan@kent.edu
Weiwen Jiang, George Mason University, wjiang8@gmu.edu
Bo Fang, Pacific Northwest National Laboratory, bo.fang@pnnl.org

Topic: Codesign and integration across the quantum computing and networking stacks

Challenge The quantum noise issue becomes more severe due to the variation of the topology (*i.e.,* connection) on physical qubits. Specifically, quantum algorithms are synthesized as quantum circuits (also called *logical circuits*). The quantum circuits will finally be deployed to the physical qubits for execution. It will go through the process of mapping logical circuits to physical qubits to obtain the physical quantum circuits (also called *compiled circuits*). The compiled circuit, however, is specified to the actual quantum computer since the topology of qubits is different. Thus, even for the same quantum algorithm, the effects of quantum noise on the algorithm can be significantly different. In NISQ era, it is crucial to improve the fidelity of quantum computing, enabling the correctness of quantum algorithms (or at least approximate to the desired results). To this end, designers of quantum algorithms need to have a clear understanding of the hidden noise of different quantum computers as well as the noise inside different compiled circuits on the same quantum computer for a specific quantum algorithm. However, this is not a trivial task. On the one hand, access to most of today's quantum computers has been provided by IT companies, like IBM, Amazon, and Microsoft, as a cloud service. To the best of our knowledge, no uniform tool revealing the noise exists to assist the designers. On the other hand, due to the rapid increase of quantum computing users, we face increasing queuing time for accessing cloud quantum computers without noise awareness. Thus, the evaluation and mitigation of noise before execution have become even more urgent to avoid the waste of queuing time for unsatisfied results. Currently, the common practice to obtain less-noisy execution results is still a trial-and-error process.

Opportunity To address the above problems, it is essential to have a tool to effectively inform users of the noise in quantum computing for a better selection of quantum computers and compiled circuits. Given that visualization has shown great power in various applications, we aim to develop a novel visualization approach to enhance users' noise awareness in quantum computing and make quantum computing more transparent and reliable. However, there exist many challenges that mainly come from two perspectives: *complex and dynamically-evolving quality of quantum computers* and *significant variations of the compiled circuits*. First, the performance of quantum computers relies on multiple factors of qubits

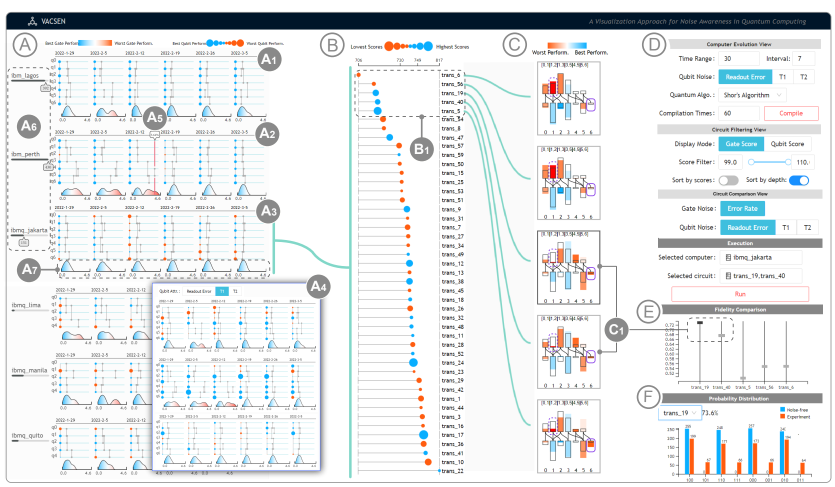

Figure 1: The interface of *VACSEN* makes users aware of the quantum noise via three linked views (A-C).

and quantum gates, such as quantum decoherence, gate error, and qubit readout error[1, 2]. These factors are dynamically changing over time [3]. It is challenging to visualize these complex factors as well as the qubit topological connections along a timeline. Second, a quantum algorithm can be compiled to various compiled circuits with significantly-different noise on the same quantum computer. For a large-scale quantum algorithm, the compiled circuits can be several hundreds. The noise of compiled physical circuit needs to be evaluated from different perspectives, *e.g.*, the circuit depths and the noise of involved qubits or quantum gates. But it is difficult to visually summarize a large number of the compiled circuits regarding the various noises, and enable users to select the most appropriate one shortly.

To fill the research gap, we propose *VACSEN*, a **V**isualization **A**pproa**C**h for noi**S**e awaren**E**ss in qua**N**tum computing. *VACSEN* can inform quantum computing users of the noise in quantum computers and compiled physical circuits, leading to a better execution result with higher fidelity. We follow a user-centered design process by working closely with five domain experts in quantum computing for over five months. A pilot study is conducted to derive the design requirements. These design requirements guide our subsequent visual designs for *VACSEN*. *VACSEN* mainly consists of three novel visualization views: Computer Evolution View, Circuit Filtering View, and Circuit Comparison View. Specifically, Computer Evolution View (Fig. 1A) facilitates the temporal noise assessment of quantum computers by a novel circuit-like design that reveals the qubit connectivity. Circuit Filtering View (Fig. 1B) supports the filtering of the compiled circuits, allowing users to pick the compiled circuits of interest. Circuit Comparison View (Fig. 1C) further enables a more detailed comparison of selected compiled circuits with a novel coupled bar chart design, facilitating the selection of an optimal compiled circuit for the final execution. To the best of our knowledge, *VACSEN* is the first visualization approach for real-time noise awareness in quantum computing.

Assessment To evaluate the usefulness and effectiveness of *VACSEN*, we present two case studies on both small-scale and large-scale quantum algorithms and conduct in-depth interviews with quantum computing users. Further assessment can be extended by integrating the visualization system into the existing Oak Ridge Leadership Computing Facility (OLCF) Quantum Computing User Program (QCUP).

Timeliness/Maturity *VACSEN* has been demonstrated in the IEEE VIS conferences in 2022[4]. The team has presented the tutorial in IEEE QuantumWeek, 2022, and Embedded Systems Week (ESWEEK), 2022. *VACSEN* is open-sourced and available to the public via `https://vacsen.github.io/`. The current timeline is: 1). system integration into the existing user program (e.g., OLCF); 2). feature enhancement based on user feedback; 3). Cross-vendor dashboard system design and implementation (e.g., IBMQ, Rigetti, Honeywell, et al).

References

[1] Abdullah Ash Saki, Mahabubul Alam, Koustubh Phalak, Aakarshitha Suresh, Rasit Onur Topaloglu, and Swaroop Ghosh. A survey and tutorial on security and resilience of quantum computing. In *2021 IEEE European Test Symposium (ETS)*, pages 1–10. IEEE, 2021.

[2] Maximilian Schlosshauer. Quantum decoherence. *Physics Reports*, 831:1–57, 2019.

[3] Betis Baheri. Temporal analysis of quantum errors in nisq computers: an empirical study. *IEEE Dataport*, pages 163–169, 2021.

[4] Shaolun Ruan, Yong Wang, Ying Mao, Weiwen Jiang, Shuai Xu, and Qiang Guan. Vacsen: A visualization approach for noise awareness in quantum computing. In *IEEE VIS*, 2022.

Towards Real-time Characterization and Control of Quantum Systems: Integrating Measurement-informed Optimal Control in a Hybrid Quantum-Classical Computing Framework[1]

Stefanie Günther, N. Anders Petersson, Sohail Reddy
Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, CA
{guenther5, petersson1, reddy6}@llnl.gov

Topic: Error mitigation, compilation, codesign and integration

Challenges: The era of Noisy Intermediate-Scale Quantum (NISQ) computers is characterized by small- to medium-scale quantum processors that are error-prone due to various effects, such as crosstalk, distortion of control signals, and environmental interactions, leading to significant errors in state preparation, quantum logical operations and measurements. Approaches to reduce these errors and enable fault-tolerant quantum computing include hardware improvements to increase qubit lifetimes, as well as the development of error correction methods. The latter, however, requires gate errors that are smaller than what current hardware can achieve, posing a significant challenge for the usability of NISQ devices. Numerical optimal control (OC), as a third pillar for achieving fault-tolerant quantum computing, has shown promising results in improving error rates of fundamental quantum operations, but it has predominantly been demonstrated for small quantum systems [1]. For larger systems, OC holds the potential to define a set of control pulses to realize entire multi-qubit algorithms, thereby realizing the desired operations directly, rather than decomposing algorithms into many single- and two-qubit gates. This could drastically reduce execution times such that high-fidelity operations can be achieved on many-qubit systems.
Unfortunately, existing methods for optimal control do not scale well with increasing number of qubits. Since the quantum state space grows exponentially with the number of qubits, simulating and optimally controlling many-qubit quantum dynamics using classical computers will eventually be out of reach, even when highly distributed high-performance computing (HPC) platforms are utilized. Further, while OC is well-developed for single and two-qubit operations where it has been successful in reducing error rates on the numerical models they are trained for, applying these controls on actual quantum devices often leads to increased error rates, as the underlying noise sources and dynamics are not fully understood, leading to discrepancies between numerical models and experimental applications [2]. As system parameters can drift over time, frequent recalibration of the model and control pulses is typically required [3]. An additional complication is that the optimization landscape often is highly non-convex and can exhibit barren plateaus when multi-qubit optimization is considered, hindering optimization progress [4].

Opportunities: For scaling up optimal control to larger quantum systems, a unified approach that integrates system characterization into the control design process is needed. In particular, data-driven approaches that inform the control pulse updates by system measurements are essential to fully exploit the potential of OC for multi-qubit operations on NISQ devices. For example, approaches such as Bayesian experimental design [5], which deal with noisy measurements using stochastic methods, can be extended for simultaneous characterization and optimal design of control pulses using measurement data. Incorporating measurement within a hybrid quantum-classical optimization is currently state of art in variational quantum algorithms (VQA), most prominently for computing eigenstates of Hamiltonian systems (VQE) [6]. However, these methods typically rely on a parameterized gate-based ansatz and their extension to pulse-level optimization is only in its nascent stages, operating on small system devices [7].

---

[1]This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-MI-848281

To enable design of optimal control pules for realizing entire algorithms at scale, it is crucial to integrate techniques for optimal control and characterization that, to a large extent, rely on the quantum device, without the need for classical simulations. New methodologies that leverage the rich toolbox of classical control theory need to be developed and generalized to incorporate measurements from the quantum device itself. For example, gradient-based methods would require efficient adjoint-based backpropagation based on measurement outcomes that can be executed on the quantum system. As an alternative, gradient-free methods should be considered, or methods that require gradients only on the classical computing side. Stochastic optimization methods, such as those used in machine learning, should be investigated systematically in the context of a hybrid quantum-classical characterization and control design process, and applied on multi-qubit systems. Further, optimization algorithms need to be developed that can operate on reduced measurements to ameliorate the high cost of state tomography. This includes researching alternate measures for large scale hybrid pulse-level optimization.

Assessment and Timeliness: Ideally, optimal quantum control techniques should be advanced to a stage where characterization and control can be done on-the-fly and in real-time. With such a capability, calibrated pulse sequences could be generated when they are needed, just before algorithm execution, such that they are informed by and aligned with the current system parameters and noise level. This ambitious goal requires close integration of the classical and quantum workflows and necessitates multidisciplinary research efforts to enable a hybrid quantum-classical computing setting. Successful application of the techniques described above can significantly improve the performance and efficiency of quantum algorithms in the NISQ era, quantified by higher fidelities of multi-qubit operations and faster algorithm executions, and could enable demonstration of quantum advantage. Furthermore, a unified data-driven approach to characterization and control can enable scientific discovery for identifying underlying physical processes in the noisy quantum computer, which in turn can lead to improvements in quantum hardware to enable longer qubit lifetimes and improved performance. In the longer term, a tight coupling between classical and quantum computing workflows would enable future quantum-enhanced technologies for scientific computing, with a wide range of practical applications (e.g. [8]).

References

[1] Koch, C. P., et al. "Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe." *EPJ Quantum Technology* 9.1 (2022): 19.

[2] Preskill, J. "Quantum Computing in the NISQ era and beyond", Quantum 2 (2018): 79.

[3] Carvalho, A., et al. "Error-Robust Quantum Logic Optimization Using a Cloud Quantum Computer Interface", *Phys. Rev. Appl.* 15 (2021): 064054.

[4] Ge, X., et al. "The optimization landscape of hybrid quantum–classical algorithms: From quantum control to NISQ applications." *Annual Reviews in Control* (2022).

[5] McMichael, R. D., et al. "Sequential Bayesian experiment design for adaptive Ramsey sequence

[6] Kandala, A., et al. "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets." *Nature* 549.7671 (2017): 242-246.

[7] De Keijzer, R., et al. "Pulse based Variational Quantum Optimal Control for hybrid quantum computing." *Quantum* 7 (2023): 908.

[8] Joseph, I., et al. "Quantum computing for fusion energy science applications." *Physics of Plasmas* 30.1 (2023): 010501.

Unsupervised Spatio-temporal Classification and Anomaly Detection using Quantum Variational Inference

Sreelekha Guggilam, Vandy Tombs, Jason Wohlgemuth

Primary Topic: Algorithms **Secondary Topics:** Applications, Computing and programming models

Challenge

Unsupervised models for spatio-temporal classification and anomaly detection have gained interest due to their applications across multiple domains that are of particular interest to Department of Energy (DOE)/ASCR (e.g., climate, nuclear proliferation, disease outbreak, etc). This is especially relevant in modeling geospatial data over time where the data spans large geographic coverage. However, modeling large spatial coverage introduces challenges such as loss of spatial correlations (especially when using hierarchical models), or increased model complexity due to changing spatial distributions and trends over time. A promising method models spatio-temporal data as dynamic graphs to identify spatial clusters and anomalies.

Current classical approaches that perform unsupervised classification (e.g., Bayesian mixtures and Monte Carlo Markov Chains (MCMC) [4, 8]) fail to extend to applications that require dynamic graphs that change at faster rates (e.g., location-based human mobility data). These classical methods can be computationally expensive which can significantly limit the data size that can be modeled. And those that do provide significant improvements in speed and memory fail to perform in presence of mixture distributed data (which is often seen in most spatial datasets). Although some methods like variational inference (VI) [5] may provides significant edge to classical MCMC approaches, modeling large dynamic graphs still remains a challenge and methods are needed to scale these classical methods to larger application domains. Utilizing quantum computing and quantum probability theory in conjunction with or inplace of these classical methods show a promising direction for unsupervised classification of not just spatio-temporal data, but any application domain with extensive amounts of data.

Objectives

We propose a novel unsupervised approach to model human mobility patterns, classify similar trends and detect anomalous behaviors in spatio-temporal data. A possible approach could be the use of quantum variation inference (QVI) in combination with Bayesian mixtures for classification and large deviations theory for anomaly detection. For this, the spatio-temporal data must be represented as dynamic graphs. Next, estimate the transition matrices for network state transformations assuming a discrete state discrete time non-stationary Markov chain using classical variational inference as well as a hybrid quantum variation inference approach. Finally, use the transition matrices to identify underlying clusters (cliques) and anomalies (both spatial and temporal anomalies).

Ideal implementation of this approach would include the following are the four main objectives: (O1) develop the probabilistic anomaly scores for Gaussian mixtures using quantum probability theory and large deviations theory for non-temporally evolving settings, (O2) implement a classical variational inference approach for detection of changes in state space of the dynamic graphs and anomaly detection for non-stationary Markov Chains, (O3) design a novel quantum variational inference approach that extends the model from O2 to a hybrid quantum computing setting, (O4) derive metrics to quantify quantum advantage over classical machine learning models and use these metrics to compare the performance of classical and quantum variational inference models for spatiotemporal data.

Opportunity

Hybrid quantum-classical machine learning could provide methods to improve unsupervised classification and anomaly detection models for spatio-temporal data. Classical methods are generally insufficient for modeling large spatial coverage given challenges provided above. The use of quantum methods like quantum variational inference approaches [1, 9, 7] applied to dynamic graphs could enable broader adoption of unsupervised classification models which may lead to extending the theory to identify extreme events in climate data, detecting or predicting machine failures in large energy grids, and better modelling of disease outbreaks. Given the immense human impact of some of these applications, it is necessary for further research towards explainability and interpretability of quantum algorithms. The research will guide the Department of Energy (DOE) and other federal agencies that prioritize spatiotemporal quantum computing in creating new scientific abilities and expanding the theory to recognize climate data's extreme events, anticipate or detect failures in extensive energy grids, and forecast disease outbreaks.

Building Theoretical Aspects of Quantum Computing: While quantum theory has shown promise in various application domains [11, 6, 3, 10, 2], there is still persistent gap between existing theory and deployable solutions. Developments in quantum probability theory are needed to effectively utilize quantum algorithms for unsupervised classification and anomaly detection models. Bridging the gap between theoretical potential and practical applications of quantum technologies requires further research and development efforts.

Quantifying of Quantum Advantage: Though multiple efforts have been made towards advancing quantum computing capabilities, development of metrics that can accurately quantify the advantage of quantum algorithms over classical machine learning approaches is still a challenge. Multiple aspects of quantum algorithms such as model complexity, possible speedup, scalability due to increased storage capabilities, the effects of decoherence, noise, and error correction etc are often expressed without any mathematical appraisal, thereby misrepresenting the algorithms' true potential. Further research and development to develop robust and standardized metrics are needed for quantifying quantum advantage across multiple domains and applications.

Assessment

The definition of success is fairly broad for this effort and could range anywhere from verifying the viability of algorithms that fail to perform better than classical approaches to validating the quantum advantage with quantifiable speed up over the analogous classical methods. That is to say, the output need not be a quantum algorithm that out performs its classical counterparts, but at least an initial understanding of the limits theoretically and physically imposed on the quantum algorithms.

Timeliness or maturity

Quantum computing and corresponding hardware has seen recent advancements and the opportunities to apply quantum computing on a variety of domains is becoming more of a reality. In the near future, quantum computing will become more accessible. Ideally, the necessary theoretical methods and algorithms will be developed along side quantum computing and hardware allowing a smooth transition from classical methods to quantum approaches. As the need for processing and making sense of the vast amounts of data continues to grow, being prepared to move quickly to quantum methods will become even more pressing. By preparing the underlying theory while quantum hardware evolves and becomes more common place, researchers and practitioners will be able guide the development of quantum computing with specific applications in mind and fully utilize the quantum advantage that will likely come.

References

[1] Marcello Benedetti, Brian Coyle, Mattia Fiorentini, Michael Lubasch, and Matthias Rosenkranz. Variational inference with a quantum computer. *Physical Review Applied*, 16(4):044057, 2021.

[2] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. 549(7671):195–202.

[3] Carlo Ciliberto, Mark Herbster, Alessandro Davide Ialongo, Massimiliano Pontil, Andrea Rocchetto, Simone Severini, and Leonard Wossnig. Quantum machine learning: a classical perspective. 474(2209):20170551. Publisher: Royal Society.

[4] Sreelekha Guggilam, Varun Chandola, and Abani Patra. Tracking clusters and anomalies in evolving data streams. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 15(2):156–178, 2022.

[5] Lauren A Hannah, David M Blei, and Warren B Powell. Dirichlet process mixtures of generalized linear models. *Journal of Machine Learning Research*, 12(6), 2011.

[6] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R. McClean. Power of data in quantum machine learning. 12(1):2631. Number: 1 Publisher: Nature Publishing Group.

[7] Hideyuki Miyahara and Yuki Sughiyama. Quantum extension of variational bayes inference. *Physical Review A*, 98(2):022330, 2018.

[8] Carl Rasmussen. The infinite gaussian mixture model. *Advances in neural information processing systems*, 12, 1999.

[9] Issei Sato, Kenichi Kurihara, Shu Tanaka, Hiroshi Nakagawa, and Seiji Miyashita. Quantum annealing for variational bayes inference. *arXiv preprint arXiv:0905.3528*, 2009.

[10] M. Schuld, I. Sinayskiy, and F. Petruccione. An introduction to quantum machine learning. 56(2):172–185.

[11] Yao Zhang and Qiang Ni. Recent advances in quantum machine learning. 2(1):e34. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/que2.34.

QUANTUM CIRCUIT COMPILATION IN A NOISY WORLD

TOBIAS HAGGE, PNNL (TOBIAS.HAGGE@PNNL.GOV)

1. Topic

This position paper addresses challenges and opportunities for quantum circuit compilers for error-mitigating and error-correcting circuits.

2. Challenge

Today's noisy quantum computers are limited to short circuits with few qubits, and hybrid (variational) algorithms. These limitations have driven the development of disparate error-mitigation techniques, including circuit randomization [3], noise injection to cancel errors in expectation values or to estimate zero-noise values [7], and error-detection and/or limited error-correction using conserved symmetries in the problem to be solved [1], or the operator algebra in which it is encoded [6, 4]. The latter methods are based on stabilizer codes and borrow techniques from quantum error-correction with an in-practice error detection capabilities which greatly exceed correction guarantees.

Quantum circuit packing and optimization techniques frequently reorder circuits using commutation relations. However, quantum operations generally lose their commutativity in the presence of noise, with the result that reordering choices can significantly affect error-mitigation performance.

It is generally agreed that for the most mature quantum technologies, significant hardware engineering challenges must be overcome to build quantum computers with more than a few thousand qubits. Until that occurs, computations will be extremely resource-constrained even if they are fault-tolerant. Compilers will need to understand how to use limited resources to maximize the probability that a computation does not stray from the error-correctable portion of the state space. In this sense, the challenges for compiling error-mitigated circuits and for compiling early fault-tolerant era circuits are the same.

3. Opportunity

There is thus an immediate, and durable, codesign need for compilers that perform reordering transformations based on their effects on the performance of the error mitigation system in the presence of hardware noise. Architectures should be scalable, and modular; developers of error-mitigation techniques should be able to leverage existing circuit optimization code without unnecessary effort. Systematic frameworks for error mitigation (such as [2]) can be built upon, and leveraged, and created to allow specification of compiler interfaces for many classes of error mitigation schemes. Stabilizer-code-based tools and algorithms can articulate the physical-space consequences of reordering transformations.

Date: May 2023.

Domain-specific languages can specify problem symmetries, which can be checked by a verifier. The action of a Clifford group element on a stabilizer code space can affect error-mitigation performance since encoded operations are intertwined with physical operations such as swaps; encodings can be optimized by searching over the action of the Clifford group on the encoded state and operator algebra.

Efficacy of error mitigation techniques can be tested beyond the bounds of physical hardware by estimating expectation values in postprocessing.[5]

4. Assessment

A successful effort will demonstrate improved accuracy, or equivalently, reduced sampling cost for a given accuracy, in a variational quantum algorithm, or else improved error characteristics in a stabilizer-code-based quantum algorithm. An even more successful effort will allow a user to specify error models and error mitigation strategies using domain languages and have a compiler produce code which is superior to code which can be produced using only one of the two inputs, with no or minimal user intervention. In a larger sense, the effort will be successful if it produces compilation tools and techniques applicable to NISQ architectures but which also find routine use in scale-limited fault-tolerant quantum computers.

5. Timeliness

The proposed work is motivated and made possible by recent expanding diversity in error mitigation techniques, work that systematizes these into frameworks, and emerging capabilities for error-correction in quantum hardware. The work will improve the efficacy of error-mitigation techniques and fault-tolerant computation. The proposed work will reduce development time and lower barriers to adoption by developing automated methods to interface error-mitigation techniques with circuit reordering techniques.

References

[1] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien. Low-cost error mitigation by symmetry verification. *arXiv e-prints*, page arXiv:1807.10050, July 2018.

[2] Zhenyu Cai. A Practical Framework for Quantum Error Mitigation. *arXiv e-prints*, page arXiv:2110.05389, October 2021.

[3] Akel Hashim, Ravi K. Naik, Alexis Morvan, Jean-Loup Ville, Bradley Mitchell, John Mark Kreikebaum, Marc Davis, Ethan Smith, Costin Iancu, Kevin P. O'Brien, Ian Hincks, Joel J. Wallman, Joseph Emerson, and Irfan Siddiqi. Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor. *Phys. Rev. X*, 11:041039, Nov 2021.

[4] Zhang Jiang, Jarrod McClean, Ryan Babbush, and Hartmut Neven. Majorana Loop Stabilizer Codes for Error Mitigation in Fermionic Quantum Simulations. *Physical Review Applied*, 12(6):064041, December 2019.

[5] Jarrod R. McClean, Zhang Jiang, Nicholas C. Rubin, Ryan Babbush, and Hartmut Neven. Decoding quantum errors with subspace expansions. *Nature Communications*, 11:636, January 2020.

[6] Kanav Setia, Sergey Bravyi, Antonio Mezzacapo, and James D. Whitfield. Superfast encodings for fermionic quantum simulation. *Physical Review Research*, 1(3):033033, October 2019.

[7] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.*, 119:180509, Nov 2017.

Integrated Quantum Workflows with Ensembles and Distributed Computation

Kathleen E. Hamilton[1], **Alberto M. Marino, Claire E. Marvinney**
Computer Science and Engineering Division, Oak Ridge National Laboratory,
Andrea Delgado, *Physics Division*, Oak Ridge National Laboratory.
Ang Li, *Physical and Computational Sciences*, Pacific Northwest National Laboratory
Shinjae Yoo, *Computational Science Initiative*, Brookhaven National Laboratory

TOPICS — Codesign and integration across the quantum computing and networking stacks, Applications, Computing and programming models

We advocate for continued investment by the Department of Energy (DOE) into the utility of distributed quantum components by supporting research into the co-design, deployment, and evaluation of quantum ensembles and quantum networks. In this paper we address both *ensembles*, i.e., a computational model composed of multiple local models, and *networks*, i.e., a system of interconnected quantum devices, as they share similar challenges in terms of scaling and optimizing construction to maximally leverage quantum effects. Distributed quantum computing offers new possibilities for both areas [1] but further research is needed to explore their full potential and address the challenges that arise in their implementations.
CHALLENGES — Designing parameterized circuits for scalable applications deployed on noisy intermediate scale quantum (NISQ) processors are faced with the barren plateau problem [2], which prevents brute force scaling of a single circuit model used in variational quantum algorithms (VQAs) (e.g. quantum neural networks). Networks that can robustly share information across nodes face similar challenges in terms of scalability and identification of optimal architectures to mitigate the effect of imperfections of current NISQ devices. Together, the use of quantum networks to connect quantum computing devices to execute distributed workflows introduces additional challenges that need to be addressed in order to take advantage of fault-tolerant quantum computing devices connected by quantum links.
OPPORTUNITIES — In general, it can be shown that the individual learners of an ensemble can converge to different local minima. As a result, ensembles constructed through bagging and boosting or geometric methods can exhibit improved performance by reducing variance or integrating information learned from different feature channels of individual learners. Ensemble models have been translated into quantum machine learning models [3], and using smaller circuits can avoid noise scaling issues in the NISQ era.

There are immediate opportunities for research into ensemble dimensionality, capacity and uncertainty quantification – understanding how noise perturbations of each learner impacts outcome robustness, or how quantum entanglement affects the dimensionality of ensembles. Additional opportunities exist in the fields of edge computing, federated learning, and privacy-preserving models which leverage ensembles to great benefit [4]. For NISQ devices, distributive asynchronous VQAs [5] have shown potential in averaging out NISQ machine-specific bias and significantly accelerating the training speed. As hardware matures, adaptive measurements also open new opportunities for the growth and construction of ensemble models. In designing distributed computing models, is there a new programming model, such as a quantum shared memory model [6], that can better connect the components of the ensemble?

Finally, can emerging collective quantum communication primitives [7] and optimizations [8] benefit ensembles? With the emergence of quantum networking and interconnects [9], distributed quantum computing provides new opportunities for quantum ensembles [10]. Implementing a quantum ensemble of interconnected devices can lead to significant advantages over independent ensembles of either classical or quantum systems. For example, a

[1]hamiltonke [at] ornl [dot] gov

quantum interface between an ensemble of quantum sensors can enable a significant sensitivity enhancement beyond that of classically connected single quantum sensors [11]. Similar advantages can be obtained for quantum interconnected quantum computers. The optimization of interconnections and readout techniques needed to better leverage quantum resources in the ensemble or network can couple to the improvement in quantum algorithms.

Combining ensemble models and networks together create the opportunity for direct learning on quantum states using edge devices that generate quantum data from quantum sensor networks, or distributed quantum computers [12, 13]. Quantum ensemble models also provide testbeds for implementing privacy-preserving federated learning at the edge.

ASSESSMENT — The needs of quantum networks and quantum computers continue to grow in size and complexity. As they scale up, the quantum state generation required for different architectures must be optimized. This is especially relevant for networked sensors to reach the optimal sensitivities and enhancements possible in the NISQ era. Achieving higher quality readout on entangled states is a metric of success for networks. A similar metric can be applied to ensemble models.

Another metric is throughput scalability– by combining measurements into a composite output using classical post-processing, or using a quantum model, is it better to scale up one single quantum network or build up a composition of smaller, distributed networks that are weakly coupled. Using smaller distributed networks can bring quantum computing closer to edge computing, however, achieving this requires robust probe state preparation. Improved processing of training features (throughput) will lead to a significant impact on advancing the state of the art in quantum applications.

MATURITY— The continued maturity of hardware make longer circuit depths and better quality sensing data possible– potential solutions will be evaluated in the quality of measurements and control precision. Improvements in hardware noise, specifically the reduction of correlated noise, make proof-of-concept quantum ensembles feasible. These models can be implemented and realized on near-term hardware and have the potential to scale with hardware sizes. As quantum networking infrastructure and quantum interconnect technology matures, there are many areas of collaboration between algorithm development and network design.

[1] Hamza Jnane et al. Multicore quantum computing. *Physical Review Applied*, 18(4):044064, 2022.

[2] Jarrod R McClean et al. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.

[3] Maria Schuld et al. Quantum ensembles of quantum classifiers. *Scientific Reports*, 8:2772, 2018.

[4] Xuan Gong et al. Ensemble attention distillation for privacy-preserving federated learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15076–15086, October 2021.

[5] Samuel Stein et al. EQC: Ensembled quantum computing for variational quantum algorithms. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, pages 59–71, 2022.

[6] Robert Beals et al. Efficient distributed quantum computing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 469(2153):20120686, 2013.

[7] Thomas Häner et al. Distributed quantum computing with QMPI. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–13, 2021.

[8] Anbang Wu et al. CollComm: Enabling efficient collective quantum communication based on EPR buffering. *arXiv:2208.06724*, 2022.

[9] James Ang et al. Architectures for multinode superconducting quantum computers. *arXiv:2212.06167*, 2022.

[10] Luca Pezzè et al. Quantum metrology with nonclassical states of atomic ensembles. *Reviews of Modern Physics*, 90:035005, 2018.

[11] Benjamin J Lawrie et al. Quantum sensing with squeezed light. *ACS Photonics*, 6:1307, 2019.

[12] C Huerta Alderete et al. Inference-based quantum sensing. *Physical Review Letters*, 129(19):190501, 2022.

[13] Yi Xia et al. Quantum-enhanced data classification with a variational entangled sensor network. *Physical Review X*, 11:021047, Jun 2021.

Towards High-Level Quantum Programming Languages

Daniel Huang* SFSU
E Wes Bethel SFSU
Talita Perciano LBNL
Roel Van Beeumen LBNL
Daan Camps LBNL

Topic: Computing and programming models: (1) design and analysis of established and novel abstract quantum computing and programming models and (2) models for hybrid quantum and classical computing.

1 Challenge

Throughout the history of classical computing, programming languages have been used to simplify the process of programming so that we can more easily take advantage of advancements in hardware. However, existing quantum algorithms are still implemented in low-level languages (e.g., hand-crafted circuits or Hamiltonians). This makes it difficult to develop new algorithms, compare algorithms across different models of quantum computing (e.g., circuit versus adiabatic), and handle limitations of noisy intermediate-scale quantum (NISQ) devices. This raises the question: can we design a high-level quantum language? This would raise the level of abstraction, thereby enabling code reuse, modularity, platform independence, and platform heterogeneity that is found in the classical setting.

Unfortunately, it is difficult to develop a "simple" model of a high-level quantum programming language to begin answering this question. In particular, such a high-level language will necessarily involve interaction between quantum, classical, and probabilistic computation due to *measurement* and *entanglement*. To see this, consider the following program `partial_measure_bell` that partially measures a Bell state.

```
1 def partial_measure_bell():
2   x = 1/√2(|00⟩ + |11⟩)
3   y = measure(x[0])
4   return y, x[1]
```

On Line 2, we construct a Bell state and store the result in variable `x`. On Line 3, we measure qubit 0 of `x` and store the result in variable `y`. On Line 4, we return both the result of measurement `y` and the unmeasured qubit `x[1]`. The unmeasured qubit `x[1]` is a mixed state that takes on $|0\rangle$ with probability 0.5 and $|1\rangle$ with probability 0.5. Even in this example, there are already numerous semantic issues.

First, how does quantum state (variable `x`) and classical state (variable `y`) in the language interact? In particular, `x[1]` is entangled with `y` so that if the value of `y` is known, then the value of `x[1]` will be as well. Thus we have questions related to hybrid quantum and classical computation. We also have issues of heterogeneous compute where multiple kinds of hardware devices are programmed arise. It may, for example, take longer to access the quantum state than the classical state depending on the architecture. This has implications for practically efficient algorithm design.

Second, what model of quantum computation do we want to expose at the higher level language? Many existing quantum languages today focus on a circuit-model of computation. However, there exist other promising models of quantum computing such as those based on adiabatic quantum computation [2]. To compare with the classical setting, programming languages often blend aspects of different models of classical computing. For example, imperative programming languages, such as `C`, provide language features for (1) addressing memory (abstraction of Turing machine tape) and (2) function pointers (restricted form of higher-order functions from lambda calculus). Functional programming languages, such as `Scheme`, provide language features such as (1) higher-order functions that are inspired by lambda calculus and (2) references (abstraction of Turing machine tape). What would such a quantum language that blends aspects of different models of quantum and classical computing look like?

Third, how does probabilistic computation interact with quantum state? Notably, the variable `y` is a random variable since it takes on value of $|0\rangle$ or $|1\rangle$ with probability 0.5. Thus the result of a computation is probabilistic, and values of a more complex computation may need to be inferred from multiple runs. Additionally, the noisy nature of NISQ hardware may increase the difficulty of drawing inferences from the results of computation.

*email:danehuang@sfsu.edu

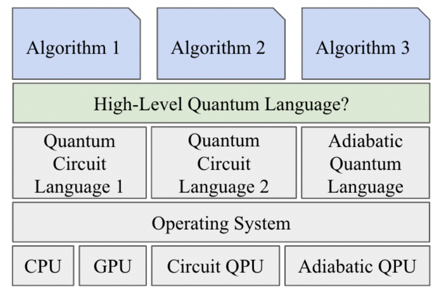

Figure 1: Opportunity to develop high-level quantum programming languages for expressing algorithms that leverages lower-level quantum software stack.

2 Opportunity

We highlight the opportunity for the development of high-level quantum programming languages (Figure 1). The high-level quantum programming language sits upon a lower-level quantum software stack comprised of lower-level languages that provide an interface between software and hardware. Such lower-level languages perform important optimizations, such as circuit compression and qubit mapping. Our concern is how to leverage this lower-level stack to develop higher-level languages for raising the level of abstraction and enabling novel applications. We believe that it will be an interdisciplinary effort to develop a high-level quantum programming language.

Challenge 1: Heterogeneous compute and opportunities in high-performance computing The high-performance computing (HPC) community has extensive experience with programming heterogeneous architectures. There are opportunities to apply and grow this body of knowledge to include programming quantum hardware. This would achieve a hybrid quantum and classical programming model. One example that comes to mind is adapting partitioned global address space (PGAS) languages from the classical setting. In this setting, the higher-level language helps users program a system where memory accesses have different costs depending on its locality. In the quantum setting, we might imagine a similar model for accessing qubits and classical memory along with the associated costs of accessing quantum/classical memory. Another example that comes to mind is programming a quantum device as we would a hardware accelerator such as a GPU. Again, there are questions concerning streaming and asynchronous computation between CPU and accelerator that are applicable to the hybrid quantum and classical setting.

Challenge 2: Language design and opportunities in programming languages There are opportunities to apply traditional language design to quantum languages to bridge the gap between different models of quantum computing such as circuit-based and adiabatic quantum computing. This would require the design of platform-portable intermediate representations and compilers to support targeting different quantum devices that implement various models of quantum computing. In combination with the constraints of hybrid quantum and classical, and heterogeneous programming, there is a wide space of possible language designs to explore.

The study of programming language semantics can also help us disentangle (pun intended) the interactions between quantum, classical, and probabilistic computation. For example, quantum computations are reversible, but classical and probabilistic computations in general are not. Clear semantics and clever implementation strategies such as uncomputation [4] can be used to ensure that quantum computations are reversible, while also supporting classical and probabilistic computations.

Challenge 3: Probabilistic computation and opportunities for probabilistic inference Quantum algorithms additionally require probabilistic inference on observed outputs. Tools from probabilistic inference such as Bayesian inference, which provide a principled methodology for drawing inferences from data based on probability, may be critical to inferring the results of computation. Additionally, the probabilistic framework is also helpful for modeling the presence of noise found in NISQ hardware. The use of programming language technology to automate Bayesian inference is studied in the field of probabilistic programming [12]. Such technology may be fruitful to apply in developing high-level quantum languages since inference is required anyway.

Taken together, we believe there is strong opportunity for studying high-level quantum languages that provides a unified abstraction for (1) enabling hybrid quantum and classical/heterogeneous compute, (2) blending multiple models of quantum computing, and (3) handling probabilistic and inferential aspects of the computation.

3 Assessment

There are several ways in which we can assess the design and implementation of high-level quantum languages. First, we can check that the language is semantically well-defined. This can be done by developing a formal semantics, i.e., mathematical model of a language. This can be used to guide sound implementation of languages.

Second, we can qualitatively gauge the ease of use of the language, similar to how classical langauges are evaluated. For instance, we can compare the amount of code needed to express an algorithm in a high-level language versus in current practice. It should also be possible to compare and contrast the performance of code written in a high-level language versus a lower-level language to study the trade offs.

Third, we can see how many new use cases are enabled by providing a higher-level language. One promising direction to explore includes studying how such a language can be used for quantum data representation and processing applied to different types of data such as images. Another direction includes domain-specific use cases such as the Variational Quantum Eigensolver [6] for solving problems in quantum chemistry which is a promising near-term application for quantum computers. Graphics and quantum ray-tracing [8] which also has applications in fusion science [5] offers another potential use case. As a final example, we may find uses cases for the high-level language in quantum machine learning [10], including quantum neural networks [1] and quantum Monte Carlo [7]. Such use cases might benefit from the high-level combination of probabilistic inference and quantum computation.

4 Timelines or maturity

There are two primary reasons why we believe now is the time to develop high-level quantum languages. First, cloud access to NISQ hardware such as IBM Quantum (`https://quantum-computing.ibm.com/services/resources`), Quantinuum (`https://www.quantinuum.com/hardware`), D-Wave's Leap (`https://cloud.dwavesys.com/leap`) are now available. Consequently, it is now possible to program such devices. Since noise, limited qubit count, and limited qubit connectivity are still issues, we will likely need heterogeneous compute to take advantage of existing hardware, which strengthens the argument for developing higher-level languages that potentially mix in probabilistic inference.

Second, there is now an emerging low-level quantum software stack providing a layer of abstraction between hardware and programming model. This stack provides circuit synthesis [13], low-level languages [3, 9, 14], and verification tools [11]. A high-level programming language can now take advantage of this software stack to continue to push the boundaries of what is possible with quantum computation.

References

[1] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner. The power of quantum neural networks. *Nat. Comput. Sci.*, 1(6):403–409, 2021.

[2] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev. Adiabatic quantum computation is equivalent to standard quantum computation. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 42–51, 2004.

[3] G. Q. AI. Cirq, 2019.

[4] B. Bichsel, M. Baader, T. Gehr, and M. Vechev. Silq: A high-level quantum language with safe uncomputation and intuitive semantics. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 286–300, 2020.

[5] A. Colaïtis, R. Follett, J. Palastro, I. Igumenschev, and V. Goncharov. Adaptive inverse ray-tracing for accurate and efficient modeling of cross beam energy transfer in hydrodynamics simulations. *Physics of Plasmas*, 26(7):072706, 2019.

[6] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *nature*, 549(7671):242–246, 2017.

[7] D. Layden, G. Mazzola, R. V. Mishmash, M. Motta, P. Wocjan, J.-S. Kim, and S. Sheldon. Quantum-enhanced markov chain monte carlo. *arXiv preprint arXiv:2203.12497*, 2022.

[8] X. Lu and H. Lin. A framework for quantum ray tracing. *arXiv preprint arXiv:2203.15451*, 2022.

[9] I. Quantum. Qiskit 0.42.1, 2023.

[10] R. D. M. Simões, P. Huber, N. Meier, N. Smailov, R. M. Füchslin, and K. Stockinger. Experimental evaluation of quantum machine learning algorithms. *IEEE Access*, 11:6197–6208, 2023.

[11] R. Tao, Y. Shi, J. Yao, X. Li, A. Javadi-Abhari, A. W. Cross, F. T. Chong, and R. Gu. Giallar: push-button verification for the qiskit quantum compiler. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, pages 641–656, 2022.

[12] J.-W. van de Meent, B. Paige, H. Yang, and F. Wood. An introduction to probabilistic programming. *arXiv preprint arXiv:1809.10756*, 2018.

[13] A. Wu and et al. A synthesis framework for stitching surface code with superconducting quantum devices. ISCA '22, 2022. `https://doi.org/10.1145/3470496.3527381`.

[14] E. Younis, C. C. Iancu, W. Lavrijsen, M. Davis, E. Smith, and USDOE. Berkeley quantum synthesis toolkit (bqskit) v1, 4 2021.

Envisioning a Middleware for Hybrid Quantum-HPC Applications

Andre Luckow[1] Nishant Saurabh[2], Pradeep Mantha[1], Shantenu Jha[3,4]

[1] *Ludwig Maximilian University Munich, Germany,* [1] *{name}.{surname}@ifi.lmu.de*
[2] *Utrecht University, NL,* [2] *{firstinitial}.{surname}@uu.nl*
[3] *Brookhaven National Lab, NY, US,* [3] *{name}@bnl.gov*
[4] *Rutgers University, NJ, US,* [4] *{name}.{surname}@rutgers.edu*

Introduction and Motivation: Quantum computing promises to accelerate scientific applications by utilizing accelerated simulations, machine learning, and optimization techniques. To support such applications and techniques, quantum computing systems are evolving from monolithic systems towards modular and distributed architectures comprising multiple quantum processing units (QPUs) coupled to classical computing nodes (HPC) (1; 2; 3). Algorithms for both *Noisy Intermediate Scale Quantum Computers (NISQ)* and *Fault-tolerant Quantum Computer (FTQC)* require the coupling of quantum and classical systems. For example, variational algorithms (4) depend on classical optimization and quantum error correction codes and require significant classical computation of the syndrome measurements. Hence, middleware systems that can facilitate the efficient coupling of quantum-classical computing are becoming critical to provide the necessary scale while accelerating development and deployment times.

We argue the need for a middleware to support the integration of quantum and classical components into workflows. We position our argument for the middleware based on a conceptual model developed by thoroughly analyzing 17 scientific application scenarios (5). To further strenghthen our position, in Figure 1 (left), we identify three types of integration between classical (HPC) and quantum tasks: *HPC-for-Quantum*, *Quantum-in-HPC* and *Quantum-about-HPC*. Each type has specific characteristics: HPC-for-Quantum requires interactions within the coherence time of the QPU. Quantum-in-HPC utilizes a classical task to orchestrate short-running quantum tasks and requires medium coupling. Quantum-about-HPC connects composable, loosely-coupled tasks to workflows.

Challenges: *Integration Patterns (challenge 1):* While quantum computers can encode any function as a classical computer, running quantum applications and workflows in practice will involve both classical and quantum tasks. Typically, only a minimal kernel, providing a quantum advantage, will often be executed on a QPU. This kernel is typically augmented with significant classical components. Characterizing the aforementioned integration and coupling patterns between classical (HPC) and quantum tasks on the middleware level is critical.

Workload and Resource Management (challenge 2): Managing quantum and classical resources can be difficult due to the varying and unpredictable resource demands, requiring a sophisticated approach to resource management. For example, for variational circuits, the QPU resource demands can vary significantly as using different optimizers, e. g., can result in a different number of circuit executions. Gradient-based optimizers require more executions of a quantum circuit to estimate the gradient using the parameter shift rule than non-gradient-based optimizers. Thus, a middleware system that can adaptively manage the resources is required. With increasing scale, data and computational requirements will become even more demanding and, thus, require efficient resource management.

Heterogeneity (challenge 3): Various quantum software frameworks have emerged (6), e. g., Pennylane (7), Qiskit (8), and Cirq (9). These frameworks support creating and executing quantum circuits on multiple quantum backends (e. g., simulators and real quantum devices) and enable interfacing with various hardware platforms (e. g., superconducting and ion trap platforms). Most frameworks provide some parallelization and accelerator support, e. g., for just-in-time compilation for optimizing circuits and supporting accelerated GPU simulators. However, they are typically limited to specific hardware platforms and do not interface with HPC resource managers. As workload and task management are deeply integrated into these frameworks, the degree of integration with HPC systems and, thus, the scale is limited.

Opportunity: A conceptual quantum middleware (Figure 1 (right)) designed based on an in-depth understanding of applications and integration patterns will accelerate the development of middleware systems (challenge 1). It will enable researchers to reason about performance trade-offs, thus enabling better and more scalable applications. Doing so will directly help to advance research into quantum algorithms, e. g., by supporting the development of new algorithms and applications that can leverage the unique capabilities of quantum computing while utilizing the existing strengths of classical computing systems.

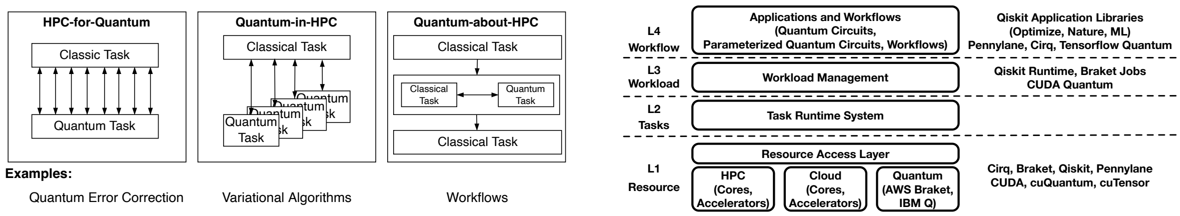

Figure 1: **Quantum-HPC Workload Patterns and Conceptual Quantum HPC Middleware:** We identify three integration patterns based on different characteristics of the quantum-classical interactions. Our conceptual middleware (5) separates and abstracts different concerns into four layers, simplifying algorithm development and deployment.

By developing a middleware that supports the integration of quantum and classical components into workflows, researchers can effectively utilize resources and tasks across quantum and classical systems, providing the critical scale to experiments using both simulated and real quantum hardware. Such middleware needs to be modular, composable, and adaptable to various application requirements and hardware configurations, thus enabling seamless integration of quantum and classical resources (challenges 2 and 3).

Assessment: Success of the envisioned middleware can be evaluated through a combination of metrics related to efficiency, scalability, and usability: (i) Increased scale and throughput: The middleware should enable the execution of larger and more complex quantum experiments, leading to better performance in terms of runtime, resource utilization, and accuracy. It should also enable the effective use of available qubits, allowing researchers to solve large and more complex scientific problems. (ii) Reduced complexity: The middleware should simplify integrating quantum and classical components into workflows, making it easier for developers to design and implement hybrid quantum-classical applications. (iii) New algorithms and improved quality: The middleware should facilitate the development of novel algorithms and optimization techniques that can provide higher quality results at a lower cost than existing methods. (iv) Adoption of the conceptual architecture: If multiple teams or organizations implement the proposed conceptual architecture in their systems, it will validate its effectiveness and utility.

Timeliness: The recent breakthroughs that make progress possible now are the advancements in quantum computing hardware. As hardware becomes more mature and available within HPC centers and cloud platforms, many applications can now experiment with quantum computing for accelerations. These developments increased the need for a middleware that seamlessly integrates quantum and classical components into workflows. As the field of quantum computing is evolving rapidly, the middleware will accelerate the ability to integrate new systems as they emerge.

References

[1] T. S. Humble, A. McCaskey, D. I. Lyakh, M. Gowrishankar, A. Frisch, and T. Monz, "Quantum computers for high-performance computing," *IEEE Micro*, vol. 41, no. 5, pp. 15–23, 2021.

[2] M. Schulz, M. Ruefenacht, D. Kranzlmüller, and L. B. Schulz, "Accelerating hpc with quantum computing: It is a software challenge too," *Computing in Science & Engineering*, vol. 24, no. 4, pp. 60–64, 2022.

[3] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, "The future of quantum computing with superconducting qubits," *Journal of Applied Physics*, vol. 132, no. 16, p. 160902, oct 2022. [Online]. Available: https://doi.org/10.1063%2F5.0082975

[4] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, "Variational quantum algorithms," *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021. [Online]. Available: https://doi.org/10.1038/s42254-021-00348-9

[5] N. Saurabh, S. Jha, and A. Luckow, "A conceptual architecture for a quantum-hpc middleware," in *submission to 2023 IEEE International Conference on Quantum Software*, 2023.

[6] M. A. Serrano, J. A. Cruz-Lemus, R. Perez-Castillo, and M. Piattini, "Quantum software components and platforms: Overview and quality assessment," *ACM Comput. Surv.*, vol. 55, no. 8, dec 2022. [Online]. Available: https://doi.org/10.1145/3548679

[7] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, J. M. Arrazola, C. Blank, A. Delgado, S. Jahangiri *et al.*, "Pennylane: Automatic differentiation of hybrid quantum-classical computations," *arXiv:1811.04968*, 2018.

[8] A. Cross, "The ibm q experience and qiskit open-source quantum computing software," in *APS March meeting abstracts*, vol. 2018, 2018, pp. L58–003.

[9] A. Hancock, A. Garcia, J. Shedenhelm, J. Cowen, and C. Carey, "Cirq: A python framework for creating, editing, and invoking quantum circuits."

Toward Reliable and Perceptual Quantum Computing on the Near-Term Quantum Devices with Unstable Noise

Weiwen Jiang[1,*](wjiang8@gmu.edu), Qiang Guan[2](qguan@kent.edu)
[1] George Mason University, [2] Kent State University, [*] Presenter

A. Topic: Error correction and mitigation; Compilation; Co-design and integration.

B. Challenge: The inherent noise, due to the sensitivity of the quantum system to the environment, sets up barriers for the practical use of quantum devices. On the one hand, the result of a quantum circuit can be far different from the theoretical or simulated result due to the influence of noise, and it essentially requires a `reliable` system design to be resilient to noise. On the other hand, quantum users can only see calibrated noise data from vendors, but they can hardly infer the performance of the application, and it is crucial to have an awareness of quantum noise on system performance, called `perceptual` quantum computing.

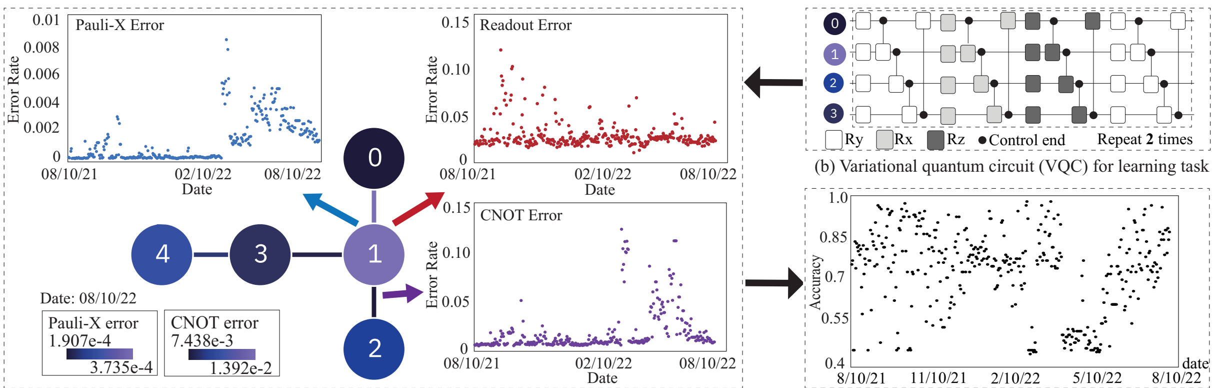

(a) Fluctuating quantum noise on real quantum computer (1-year long daily profiling)
(b) Variational quantum circuit (VQC) for learning task
(c) Fluctuating noise leads to varied accuracy of one VQC

Figure 1: Results of one type of quantum application (i.e., variational quantum circuit VQC-based quantum learning) on IBM quantum processor with fluctuating noise: (a) noise data (gate error, readout error) of *ibmq_belem*; (b) VQC circuit; (c) unstable prediction accuracy caused by noise.

Recently, our preliminary works [1, 2, 3] have shown that the near-term quantum devices will have not only high noise but also unstable noise. That is, the noise is dynamically changing over time. This can be observed from a 1-year profiling on an *ibmq_belem* quantum computer shown in Figure 1(a), where the error rates of different error sources (e.g., gate error [4]) change along with time, called unstable noise.

The unstable noise will make the design of a reliable quantum system more challenging; meanwhile, it makes perceptual quantum computing more important. We created a Variational Quantum Circuit (VQC) with 4 qubits, 80 parameters, and a length of 50, as shown in Figure 1(b), for an earthquake detection task. It obtains 98.0% prediction accuracy on perfect simulation to process 1,500 seismic data, which is even higher than the classical counterpart with the same number of parameters (i.e., 75.3%). However, by executing our pre-trained VQC under the unstable noise in Figure 1(a), we obtain the detection accuracy in Figure 1(c), where we can easily see the accuracy fluctuate heavily and the lowest accuracy is merely 44.7%.

With the property of unstable noise, reliable and perceptual quantum computing needs to be redefined. `Temporal Reliability:` The system needs to be resilient to different quantum noise in terms of time for stable and high performance. `Continuous Perception:` The assessment of the performance of a quantum circuit on a quantum device has to be frequently conducted along with the change of noise. The new requirements bring several new challenges: `(1) Lightweight Noise Adaptation:` The ideal solution for achieving temporal reliability is to find a robust circuit that can be resilient to all noise; however, this is almost impossible since the noise changes in a random fashion. Alternatively, a must-do task is to adapt a solution to the current noise. Since the adaptation is performed at run-time, efficiency is of utmost importance. Therefore, innovation for efficient noise adaptation is highly demanded. `(2) Real-time Uncertainty Quantification:` Benchmarking quantum gates for their noise has a high cost (30-90 minutes for IBM quantum computers); not to mention quantifying the application performance for an entire quantum circuit. What's worse, the uncertainty quantification needs to be frequently conducted since users need to understand whether it is worthwhile to execute a quantum circuit on the quantum device at the current noise state; what's more, a lengthy performance probe may lead to the assessment being invalid since the noise may have already changed substantially again.

C. Opportunity: The `noise adaptation` can be achieved via optimizations at different layers: (1) At the quantum circuit layer, our recent work [2] show that the performance can be stabilized by adjusting the parameters and compressing the quantum circuit; however, it has scalability issue when the compression performed on a large-size circuit. Therefore, new optimization approaches that can improve scalability need to be developed; for example, we can integrate the layer freezing technique in the optimization. (2) At the pulse control layer, a run-time calibration to adjust the amplitude and duration can provide precise performance estimation for quantum gates. On top of this, a matched optimization framework for a dual-objective on single gate fidelity and overall circuit duration can maximize the performance. As a whole, separate optimization at different layers may not provide the best performance, and a co-design is essential to further improve the performance. For `uncertainty quantification`, our preliminary work [5] shows that Machine Learning (ML) can be used; however, it is still challenging to make the quantification process to be real-time; in particular, for large-size circuits. Therefore, new lightweight ML approaches are needed to support real-time uncertainty quantification.

D. Assessment: First, the performance and the stability of a quantum system can assess the noise adaptation. The performance metric can be *either* well-defined fidelity that measures how close the final quantum state of the real-life qubits is to the ideal case *or* an applications-specific metric, such as prediction accuracy for classification tasks in quantum learning. Then, the system stability can be assessed by obtaining the performance metric at different times. Second, in uncertainty quantification, the quality can be assessed by comparing the prediction results and results obtained on actual quantum devices. Different metrics can be used, such as Huber loss, mean square error, and mean absolute error. In addition, in the above two processes, based on the optimization approach to be used (e.g., run-time calibration), we can calculate estimates of the time needed to use quantum devices (i.e., the benchmarking of gates in calibration). It would be useful for planning the time needed for job reservations.

E. Timeliness or Maturity: We are now witnessing the scale-up of quantum systems in terms of qubits' numbers, which are expected to unleash the power of quantum computing in real applications; however, the newly identified unstable quantum noise by the team [1, 2, 3] is a major roadblock. Although a fault-tolerant quantum computer via quantum error correction is the ultimate goal, it has an extremely high cost that requires qubits numbers far beyond today's quantum computers. A more swift and agile system-level cross-layer approach can adapt quantum systems to noise at run-time, which will enable temporal reliable quantum computing. What's more, uncertainty quantification can help to better arrange the limited quantum resources for applications that can adapt to the current noise.

References

[1] Samudra Dasgupta and Travis S Humble. Assessing the stability of noisy quantum computation. In *Quantum Communications and Quantum Imaging XX*, volume 12238, pages 44–49. SPIE, 2022.

[2] Zhirui Hu, Youzuo Lin, Qiang Guan, and Weiwen Jiang. Battle against fluctuating quantum noise: Compression-aided framework to enable robust quantum neural network. *arXiv preprint arXiv:2304.04666 (Accepted by the 60th Design Automation Conference (DAC))*, 2023.

[3] Betis Baheri, Daniel Chen, Bo Fang, Samuel A. Stein, Vipin Chaudhary, Ying Mao, Shuai Xu, Ang Li, and Qiang Guan. TQEA: temporal quantum error analysis. In *51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2021, Taipei, Taiwan, June 21-24, 2021 - Supplemental Volume*, pages 65–67. IEEE, 2021.

[4] Konstantinos Georgopoulos, Clive Emary, and Paolo Zuliani. Modeling and simulating the noisy behavior of near-term quantum computers. *Physical Review A*, 104(6):062432, 2021.

[5] Hanrui Wang, Pengyu Liu, Jinglei Cheng, Zhiding Liang, Jiaqi Gu, Zirui Li, Yongshan Ding, Weiwen Jiang, Yiyu Shi, Xuehai Qian, et al. Quest: Graph transformer for quantum circuit reliability estimation. *arXiv preprint arXiv:2210.16724*, 2022.

Revolutionary Paradigms for Integrated Quantum Learning and Simulation

Author: Zhenhua Jiang, University of Dayton Research Institute (UDRI)
Email: zjiang4@uayton.edu
Topic: Quantum Algorithms- quantum algorithms admitting theoretical or empirical evidence of advantage for fundamental domains such as simulation, optimization, or machine learning

Challenge: The current practice of machine learning (ML) aims to search nonlinear functions between the input and output variables to fit the training data samples and update the weights and biases iteratively, by using the gradients calculated from the errors between the predicted and labelled values [1]-[12]. Generally, this process targets seeking unknown or hidden correlations / patterns from training data in an implicit way. That being said, *the current ML paradigm has many limitations* [6]-[12]. For instance, the iterative updating of weights is fragile, relies on the gradients, and may not accurately reflect the general correlations existing in the datasets [6]. Typically, the classical training process is time-consuming, and may sometimes result in over-fitting or poor performance [7]. Point-to-point predictions using fixed values of the learned weights may not well capture the variations or uncertainty in model parameters [8]. Slight changes in the inputs may also cause large deviations in the predictions or even wrong outputs [9]. Also, the effectiveness of traditional ML on learning physical dynamics and optimizing closed-loop control schemes are also limited [10]-[11]. Some probabilistic learning schemes exist but these solutions typically assume Gaussian distributions [12], so the expected accuracy of the learned probabilistic weights may not be guaranteed. Classical solutions to end-to-end, physics-aware, non-parametric probabilistic learning are not known to exist yet.

Recently, quantum machine learning (QML) has attracted a lot of interest, with some evidence showing advantage over traditional machine learning [11], but most of the existing QML methods are just quantum implementations of traditional machine learning algorithms, in particular, based on iterative weight updating schemes. No fundamentally-new learning paradigms exist yet that are particularly well suited to quantum implementation with explicit quantum advantages. Also, quantum machine learning may see fundamental limits. For instance, a current constraint is that the targeted problem size may be subject to the available quantum resources (e.g., number of qubits and circuit depth). Also, the quantum learning acceleration would be limited by the speed to encode classical data into quantum states, which is common to many variational quantum algorithms or solutions. In addition, the errors, noise or de-coherence appearing in the quantum system may limit scaling-up of the solution. In the noisy intermediate-scale quantum (NISQ) era, hybrid classical/quantum algorithms may be expected to demonstrate more favorable behaviors than pure classical or quantum systems in solving computationally complex problems.

Opportunity: New paradigms of quantum machine learning are urgently needed that aim to address all of these limitations and outperform the classical solutions. In particular, there is a need for innovative quantum probabilistic learning paradigms fundamentally different from traditional ML, *without* the needs for iterative gradient and weight updating or probability distribution pre-assumptions.

A unique opportunity exists for *Quantum Probabilistic Learning Machine* (QuProbLeM) – a revolutionary machine learning (ML) paradigm inspired by the philosophy of *Parts and the Total* (e.g., each part, or data sample, in a closed system affects the total in such a way that the whole system will thereafter influence every single part). This first-of-the-type learning machine can be synergistically implemented on a quantum computing system that can also capture the same philosophy naturally. This QuProbLeM solver can take as its input a set of encoded data samples comprising pairs of cause-effect variables, and its output may include the updated (learned) probability distributions of the model's parameters (i.e., weights) and the predicted effect variables and their distributions when given new cause variables or their distributions. Alternatively, classical data can be first fed into a front-end Bayesian learning network to initially learn the probability distributions of the cause-effect variable pairs and then encoded into coherent quantum states as the input. The QuProbLeM solver can directly compute the probability distributions of the learning network's weights in a superpositional, inference-based manner, but not iteratively as in traditional ML.

Another excellent opportunity exists for quantum computing in an *integrated multi-scale physics-guided learning and simulation framework* for stochastic modeling. This framework will account for dynamics of multi-scale processes and systems, from quantum levels to grid scales, and capture three aspects of data-driven modeling, including spatiotemporal dynamics, causal relationships and uncertainty quantification. The computational framework can (1) integrate data-driven machine learning capabilities with atom-level physics-based simulation and stochastic process modeling of power grids, (2) leverage quantum computing to both simulate high-fidelity physical models and learn insight or physical laws from experimental data, and

(3) implement both the learning and computing tasks on hybrid quantum/classical computing resources. The hierarchical learning mechanisms enables the low-level parameter learning to feed the higher-level dynamics learning, and the meta-learning mechanisms allows the system-level rules to guide the learning at a lower level. This unified framework can serve as a common interface for integrated learning and simulation of energy materials, devices and grids, to which existing tools may be integrated as plug-in tools.

Assessment: The novel quantum learning machines should be (1) *explainable* and applicable to different physical systems or datasets, (2) *scalable* in the training data size, allowing for life-time learning and extremely sparse datasets (i.e., few shots) and covering two extreme ends of ML, (3) *adaptable* to data novelty, capturing the evolving uncertainty in the probabilistic model parameters, and (4) *ultrafast*, due to superposition, entanglement and disentanglement in quantum computing.

We should develop benchmark algorithms and software for a generalized quantum learning problem and define metrics to compare with experimental data in terms of accuracy, computational time, resource consumption, algorithm complexity, or other metrics. A novel feature of the QuProbLeM solver is a model fidelity or likelihood metric that is defined by comparing the model prediction with measurement data based on Bayesian inference. This metric can reflect the probability of model matching and be used as likelihood or evidence to affect the activation through another round of Bayesian inference and then update the weight's probability distributions. These performance metrics should be investigated for different uncertainty sources and modeling mechanisms and validated against experiment data as well. Also we need to explore specific applications of quantum learning with a practical quantum system (expected to be available in next few years) as a co-processor for a classical computing system, to solve classically challenging problems.

Timeliness or maturity: Due to the powerful parallel processing and probabilistic computing capabilities, quantum computers can make progress feasible now, where it wasn't possible before. The new quantum learning machines may become available due to recent breakthrough in a data-driven, physics-informed, Bayesian neural learning network (or method) with probabilistic weights for trainable parameters, which can account for prior knowledge, operation data (e.g., evidence) and uncertainty in the datasets [13]-[15]. In general, this learning network can learn physical system dynamics or temporal dependence from operation data, or model inherent causal relationships between cause and effect variables. It can take deterministic datasets as input and encode them in the pure states of qubits, or accept stochastic data (i.e. probability distributions of cause/effect variables) and encode them in the mixed or pure states. The solver is working in the probability space, one dimension beyond the numerical space for classical computing, doubling the degree of freedom for computing. Without iterations, this learning machine enables a very high computing speed, and essentially may be as fast as the training data can be generated and encoded.

Impact: The new paradigms will be disruptive in the general field of machine learning, benefiting a large population of ML end users in almost every section of scientific and engineering domains in the near future. If successful, this work will have transformational impact on new materials discovery, energy conversion, renewable energy, energy efficiency, carbon reduction, climate studies, portfolio optimization, power grid planning, operation and protection, infrastructure resiliency analysis, and national security in general.

References:

[1] LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. Nature 521, 436–444 (2015). https://doi.org/10.1038/nature14539
[2] Pan, J. Machine learning proved efficient. Nature Computational Science 2, 619 (2022). https://doi.org/10.1038/s43588-022-00344-8
[3] Caro, M.C., Huang, HY., Cerezo, M. et al. Generalization in quantum machine learning from few training data. Nature Communications 13, 4919 (2022). https://doi.org/10.1038/s41467-022-32550-3.
[4] Schetakis, N., Aghamalyan, D., Griffin, P. et al. Review of some existing QML frameworks and novel hybrid classical–quantum neural networks realising binary classification for the noisy datasets. Sci Rep 12, 11927 (2022). https://doi.org/10.1038/s41598-022-14876-6.
[5] Cerezo, M., Verdon, G., Huang, HY. et al. Challenges and opportunities in quantum machine learning. Nat Comput Sci 2, 567–576 (2022). https://doi.org/10.1038/s43588-022-00311-3
[6] Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell 1, 206–215 (2019). https://doi.org/10.1038/s42256-019-0048-x
[7] Takahashi, Y., Ueki, M., Tamiya, G. et al. Machine learning for effectively avoiding overfitting is a crucial strategy for the genetic prediction of polygenic psychiatric phenotypes. Transl Psychiatry 10, 294 (2020). https://doi.org/10.1038/s41398-020-00957-5
[8] Ghahramani Z. Probabilistic machine learning and artificial intelligence. Nature. 2015 May 28;521(7553):452-9. doi: 10.1038/nature14541. PMID: 26017444.
[9] Schramowski, P., Stammer, W., Teso, S. et al. Making deep neural networks right for the right scientific reasons by interacting with their explanations. Nat Mach Intell 2, 476–486 (2020). https://doi.org/10.1038/s42256-020-0212-3
[10] Gebhart, V., Santagati, R., Gentile, A.A. et al. Learning quantum systems. Nat Rev Phys (2023). https://doi.org/10.1038/s42254-022-00552-1
[11] Biamonte, J., Wittek, P., Pancotti, N. et al. Quantum machine learning. Nature 549, 195–202 (2017). https://doi.org/10.1038/nature23474
[12] Sebastian, A., Pannone, A., Subbulakshmi Radhakrishnan, S. et al. Gaussian synapses for probabilistic neural networks. Nat Commun 10, 4199 (2019). https://doi.org/10.1038/s41467-019-12035-6
[13] Z. Jiang, J. Saurine, "Data-Driven, Physics-Guided Learning of Dynamic System Models", AIAA Science & Technology Forum, Jan. 2023.
[14] Z. Jiang, K. Beigh, "Bayesian Learning of Dynamic Physical System Uncertainty", AIAA Aviation Symposium, June 2022.
[15] Z. Jiang, K. Beigh, "Data-Driven Modeling of Dynamic Systems Based on Online Learning", AIAA Propulsion & Energy Forum, Aug. 2021.

Quantum Computing Applications in Smart Grid Research

Authors: Zhenhua Jiang, University of Dayton, zjiang4@udayton.edu
Zhi Zhou, Argonne National Laboratory, zzhou@anl.gov

Topic: Quantum Applications - New kinds of DOE science applications informed by quantum capabilities

Challenge: Many infrastructure systems such as power grids are highly dynamic and stochastic, due to the fluctuating renewable power generation and load demand, volatile prices, and vulnerable supply chains. Components in these networked engineering systems can fail or be subject to attack at any random time, presenting significant challenges for applications including contingency analysis and optimization based decision-making. Similar characteristics are present in other systems, including molecular motion, chemical reactions, logistics, and financial systems. To address the uncertainty and randomness issues, traditionally, statistical modeling methods and Monte Carlo simulation tools have been utilized to analyze operational data and extract useful insight. But solving large sets of probabilistic equations is computationally hard, and making decision under uncertainty is more challenging. Bayesian learning provides a probabilistic approach to capture uncertainty in decision-making by inferring about a posterior distribution from the prior knowledge of model parameter's distribution based on new data. However, quantifying the uncertainty and propagation in this type of probabilistic networks and then providing decision support, while considering the parameter probability distributions across a large operational range, are computationally intensive.

Recent advances in quantum systems have made quantum computing (QC) feasible in the near future. QC has the potential to transform computing paradigm and generate revolutionary impact on many domains due to its powerful computing capability on probability and natural parallel processing due to superposition and entanglement [1]. While quantum circuits show great promise for gate-based algorithms, many other quantum algorithms are formulated in specific models of quantum computation, such as quantum annealer (QA) [2], quantum approximate optimization algorithms (QAOA) [3], adiabatic quantum computation [4], or quantum walks [5]. Also, most gate-based quantum algorithms focus on achieving specific functions or processes, such as quantum Fourier transform (QFT) [6], quantum amplitude amplification [7], and quantum search [8]. As an integrated application, the HHL linear solver was proposed to solve systems of linear equations on a quantum circuit with great speedup [9], but the limitation is that the solver was not designed for stochastic systems. More importantly, a universal framework that can solve probabilistic problems with QC and facilitate new ways to understand/develop quantum algorithms is not available yet; and there is a necessity to create new mathematic frameworks to unify quantum algorithms.

Opportunity: To address these challenges, there exist opportunities about a holistic algorithmic framework (including methods, algorithms, and software) for quantum probabilistic solvers and decision engines (QuPSADE) [10], and their applications in solving a wide range of classically-hard practical problems. As an example application, the power grid is a complex system that involves numerous components and decision makers, all of which interact with each other over time and space. However, with increasing penetration of intermittent generation resources and variable load (such as electric vehicles), there is elevated uncertainty in bulk power systems. With progress in decarbonization and electrification, the grid has become more complicated due to more interaction with other infrastructure systems, such as gas pipelines and transportations. Considering impact from future climate changes can only make this trend even more intractable. Traditional deterministic solutions are inadequate to capture this uncertainty, and generally ignore risks imposed by stochastic resources. As a result, operation decisions are suboptimal, and fail to achieve the highest level of efficiency at the lowest costs under certain confidence levels. To make better decisions, situational awareness of uncertainty or risk is essential, as it allows for more efficient and robust utilization of all grid assets, and proactive control of grid devices. However, currently, risk metrics of individual assets are not available, and a framework for collective risk evaluation does not exist. Although some binary risk management mechanisms exist in grid operations, they do not consider fluid, granular risk assessment.

A potential concept is a holistic, probabilistic paradigm suited to risk assessment based on stochastic modeling at the asset and system levels and risk-informed predictive operation optimization at multiple time scales. At either level, probabilistic energy modeling or power flow prediction based on operation data provides a creditable basis for risk assessment with increased accuracy and confidence. Furthermore, asset risk updates and risk-based offer strategies are aggregated to augment system-level risk evaluation, generating a system risk index and facilitating portfolio management. In return, system risk valuation and risk-aware market mechanisms will enable optimizing system operation with explicit risk consideration at different time stages. While this is not feasible with current computing technologies, the QuPSADE quantum solution can enable this operational paradigm and directly lead to disruptive new technologies in grid operations (e.g., risk-based economic dispatch, financial products for open access of energy assets, etc.), and systematical tools to quantify, analyze and hedge risk, so as to embrace more emerging energy assets at lower costs with positive reliability impact.

Assessment: Solving linear equations is required in many problems such as the simulation, machine learning, and optimization in the above-mentioned solution. A classical method to solve a system of linear equations, Ax= b, is Gaussian elimination. This algorithm uses elementary row operations to produce a convenient matrix decomposition of matrix *A* and employs back substitution to determine the solution values [11]. Another algorithm that solves a linear equation system is the conjugate gradient descent method [12], which works well with positive semi-definite and sparse matrices. Implementation on parallel computing hardware such as FPGA (field programmable gate array) can further accelerate the solution [13].

A quantum solver for systems of linear equations, also called HHL solver (named after the authors: Harrow, Hassidim, and Lloyd) was proposed in 2008 as a quantum algorithm for solving linear systems and was reported to offer exponential speedup compared with classical counterparts [9]. This algorithm can estimate the result of a scalar measurement on the solution vector to a given linear system of equations. The main elements on the quantum solver include quantum state initialization, quantum phase estimation (QPE), inversion of eigenvalues, reverse QPE and quantum measurement. The HHL quantum algorithm was highly accurate, but hard to be implemented in the near term due to the required circuit depth. Reference [14] proposed a hybrid quantum-classical algorithm, called Variational Quantum Linear Solver (VQLS), for solving linear systems on near-term quantum computers. It is proved that $C \geqslant \epsilon^2/\kappa^2$, where C is the VQLS cost function and κ is the condition number of matrix A. On Rigetti's quantum computer, a VQLS was successfully implemented for a problem size of 1024×1024.

While seeking new quantum probabilistic linear (and nonlinear) solvers and quantum decision engines, we should develop benchmark algorithms and software for a generalized quantum problem and define metrics to compare with experimental data in terms of accuracy, computational time, resource consumption, algorithm complexity, or other metrics [15]-[16]. Also we need to explore specific applications of quantum algorithms with a quantum system of q*N (width*depth) > 50,000 (expected to be available within the next few years) as a co-processor for a classical computing system, to solve classically challenging problems.

Timeliness or maturity: With the development of many novel quantum encoders, decoders, comparators, and operators, it will be feasible to make progress in quantum probabilistic solvers and decision support, where it wasn't possible before. Due to this capability, many new quantum solutions may become available to solve practical problems. If successful, the proposed quantum algorithmic framework will directly have transformative effect on solving large systems of probabilistic linear, nonlinear or differential equations and making decision under uncertainty, in solving complex problems. The success of this project will eventually have revolutionary impact on many domains of scientific research, including but not limited to:

- Smart grid, autonomous driving, air/ground traffic control, supply chains, logistics, studies of cyber-physical systems, scientific experimental systems, and infrastructure vulnerability assessment,
- Probabilistic simulation/machine learning, which can be widely used in studies of chemical reactions, molecular dynamics, nuclear reactors, new materials, diseases, new drugs, or other scientific research,
- Weather forecasting, earthquake prediction, other relevant stochastic forecasting applications, climate change and environmental studies, and cosmos research,
- Financial risk analysis and investment management, by banks, insurance or credit card companies, financing organizations, and investment institutes.

References:

[1] Trabesinger, A. Quantum computing: towards reality. Nature 543, S1 (2017). https://www.nature.com/articles/543S1a.
[2] Pelofske, E., Hahn, G. & Djidjev, H.N. Parallel quantum annealing. Sci Rep 12, 4499 (2022). https://doi.org/10.1038/s41598-022-08394-8. https://www.nature.com/articles/s41598-022-08394-8.pdf?origin=ppub
[3] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, A Quantum Approximate Optimization Algorithm, arXiv:1411.4028 (2014). https://arxiv.org/abs/1411.4028
[4] Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A, Preda D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. Science. 2001 Apr 20; 292(5516):472-5. doi: 10.1126/science.1057726. PMID: 11313487.
[5] Andris Ambainis, Quantum Walks and Their Algorithmic Applications, International Journal of Quantum Information 1, 507 (2003). https://www.worldscientific.com/doi/abs/10.1142/S0219749903000383
[6] Weinstein, Yaakov S. et al. "Implementation of the quantum Fourier transform." Physical review letters 86 9 (2001): 1889-91. https://arxiv.org/abs/quant-ph/9906059
[7] Brassard, P. Høyer, M. Mosca and A. Tapp, Quantum Amplitude Amplification and Estimation, in Quantum computation and Information, Contemp. Math., pp. 53-–74 (2002). https://arxiv.org/abs/quant-ph/0005055
[8] Barends, R., Lamata, L., Kelly, J. et al. Digital quantum simulation of fermionic models with a superconducting circuit. Nature Commun 6, 7654 (2015). https://doi.org/10.1038/ncomms8654. https://www.nature.com/articles/ncomms8654
[9] W. Harrow, A. Hassidim and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, Phys. Rev. Lett. 103, 150502 (2009). https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.103.150502
[10] Z. Jiang, L. Nguyen, "Quantum Probabilistic Linear Solver: An Initial Study", 2023 IEEE Quantum Computing Engineering, under review.
[11] Althoen, Steven C.; McLaughlin, Renate (1987), "Gauss–Jordan reduction: a brief history", The American Mathematical Monthly, Mathematical Association of America, 94 (2): 130–142, doi:10.2307/2322413.
[12] Hestenes, Magnus R.; Stiefel, Eduard (December 1952). "Methods of Conjugate Gradients for Solving Linear Systems" (PDF). Journal of Research of the National Bureau of Standards. 49 (6): 409. doi:10.6028/jres.049.044.
[13] Jiang, Raziei, "Hardware Accelerated Linear System Solver", US Patent, https://uspto.report/patent/grant/10,713,332
[14] Bravo-Prieto, Carlos et al. "Variational Quantum Linear Solver: A Hybrid Algorithm for Linear Systems." Bulletin of the American Physical Society (2020): https://arxiv.org/abs/1909.05820
[15] Chen, Chih-Chieh et al. "Hybrid classical-quantum linear solver using Noisy Intermediate-Scale Quantum machines." Scientific Reports 9 (2019): https://www.nature.com/articles/s41598-019-52275-6
[16] Fernando G.S.L. Brandao and Krysta M Svore. "Quantum speed-ups for solving semidefinite programs". 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS). IEEE. 2017, pp. 415{426.

Title: Control System Benchmarking for Quantum Computers

Authors: Nick Johnson (Riverlane, nick.johnson@riverlane.com), Hari Krovi (Riverlane, hari.krovi@riverlane.com), Kevin Young (Sandia National Laboratories, kyoung@sandia.gov), Andrea Husseini (Riverlane, andrea.husseini@riverlane.com)

Topic: Co-design and integration (benchmarking quantum computing control systems), models

Challenge: Sophisticated control systems are essential to building scalable and robust quantum computers that will require millions of synchronized signals; however, comparison of control system hardware is only possible at the level of component specifications. Controller specifications differ dramatically between systems and there exist no clear metrics on controllers that can reliably predict qubit performance. There is a critical need to better understand the effect of control hardware on the performance of a quantum processor and to develop metrics that can benchmark and compare control systems.

Opportunity: To bridge the gap between classical control system specifications and quantum device performance, we propose to develop a new generation of benchmarks in three steps:

1. <u>Model the effect of classical noise on the output pulses of the control system</u>
This may be achieved by a simulation framework that describes voltage characteristics, timing resolution, system latencies and other aspects of the classical hardware that allow one to explore realistic pulse parameters in detail. Riverlane has developed a high-resolution simulation of our FPGA-based control system [1] for atomic qubits that may be used as the basis of this work.

2. <u>Create a connection between classical control system noise and quantum computing performance using a transduction model</u>
We propose to combine our control system simulator with a simulator of the underlying qubit and device physics to identify specific control features that influence qubit fidelity and overall quantum system performance. This understanding will allow the community to set a target level of computational performance for a quantum computer and determine the control system specifications needed to achieve it. The models developed will guide the improvement of current control systems. To achieve this, a device model is needed to translate the control system characteristics into a Hamiltonian or Lindblad master equation for the quantum system being manipulated. This can be used to build a detailed physics model of the qubits and their interaction dynamics, such as for a trapped-ion chain. The device model must incorporate all the relevant real-world properties of the quantum hardware, such as the electrode potentials applied to an ion trap chip. This level of complexity is necessary to capture the real-world effect on the qubit states, and is beyond what is currently achievable with standard, idealized qubit models. Combining these simulation tools with benchmarks at each layer of the technology stack will enable a type of back-propagation through the creation of a 'transduction model', possibly using tools from machine learning, to identify the pulse-level specifications of the control system needed to meet a given threshold of computational performance for applications.

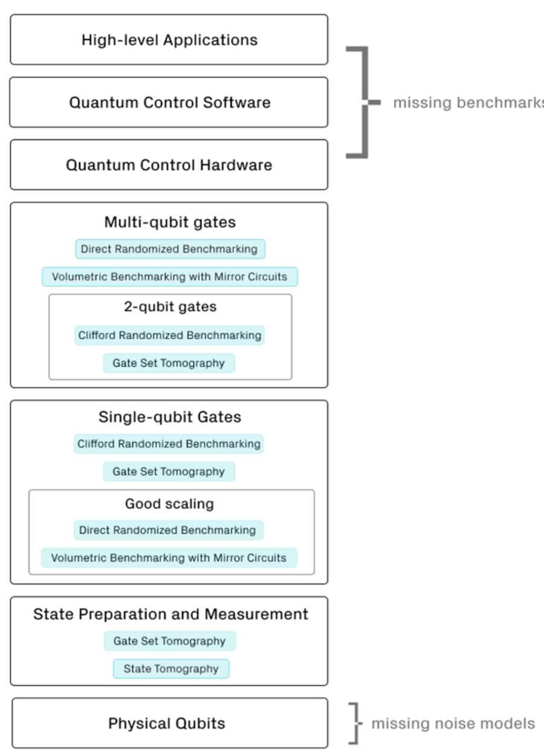

Figure 1: Missing noise models and benchmarks in the quantum stack. Established benchmarks are shown in blue.

3. Devise a suite of benchmark experiments that measure and quantify how the quality of control systems impact on qubit and system performance

These benchmarks will serve as reliable proxies for quantum behavior that can be predicted from classical measurements of controller quality. Figure 1 shows the layers of the quantum computing stack indicating where our benchmarks would fit in. Throughout the effort, we would ensure that we leverage, incorporate, and extend the existing open-source tools, such as pyGSTi [2], to enable adoption and improvement by the broader quantum computing community.

Assessment:

Successful completion of the following tasks will indicate progress on this proposal:

- Identification of specific control system features (timing accuracy, phase noise, etc.) that influence qubit fidelity and overall system performance.
- Development of a rigorously validated open-source simulator that can model the effect of control system parameters on qubit behavior and can be extended and improved.
- Creation of a database of control system features and their effect on qubit behavior in a variety of circumstances.
- Development of qubit-level and system-level benchmarks and benchmark experiments.
- Creation of a 'transduction model' so that, given a target requirement for quantum system performance, it will be possible to identify the pulse-level specifications of the control system needed to meet this requirement.
- Develop a recommended strategy to enable optimization of control systems.
- Devise a suite of experiments through which control systems can be evaluated.
- Test these experiments on a quantum computing platform.

Timeliness: Existing state-of-the-art quantum benchmarks (see, e.g., [3-9]) can only characterize the collective impact of intrinsic noise, environment errors, and controller imperfections by their effect on measured quantum circuit outcomes. These benchmarks cannot identify the *source* of errors, only their presence. The community needs methods that directly connect easily measured classical performance features of the controller to their predicted impact on qubit errors. Such tools will directly enable informed improvements and optimizations in the classical control space.

Reliable classical controller benchmarks benefit the entirety of the quantum ecosystem and will naturally lead to the development of improved control systems. Developing these benchmarks will enhance understanding of how control-level quality propagates through the hardware stack to impact device-level performance This will expedite the advancement of quantum computers, aid efforts to efficiently scale, and bring about useful quantum computing sooner.

References:

[1] Deltaflow.Control, Riverlane Ltd, www.riverlane.com/products/deltaflow-control
[2] pyGSTi, A python implementation of gate set tomography, https://www.pygsti.info/
[3] J. Eisert *et al.,* Quantum certification and benchmarking. *Nat Rev Phys* **2**, 382–390 (2020)
[4] E. Magesan *et al.,* Efficient measurement of quantum gate error by interleaved randomized benchmarking. *Phys. Rev. Lett.*, 109: 080505 (2012)
[5] E. Knill *et al*., Randomized benchmarking of quantum gates. *Phys. Rev. A*, 77:012307 (2008)
[6] J. Emerson *et al*., Symmetrized characterization of noisy quantum processes. *Science*, 317:1893–1896, (2007)
[7] T. J. Proctor *et al.,* Direct Randomized Benchmarking for Multiqubit Devices. *Phys. Rev. Lett.* 123, 030503 (2019)
[8] E. Nielsen *et al.*, Gate set Tomography. *Quantum* 5, 557 (2021)
[9] R. Blume-Kohout *et al*, A volumetric framework for quantum computer benchmarks. *Quantum* 4, 362 (2020)

Title: Quantum algorithm for data compression and preparation tools for generic Hamiltonian problems.

Authors: Pejman Jouzdani [1], Matthew Cha, Alex Haid, Mike Jaris, Mark Kostuk, Nils Paz, David Schissel, and Dan Sioui.

Affiliation: General Atomics, 3550 General Atomics Ct, San Diego, CA 92121, USA.

Topic: Quantum Algorithms.

CHALLENGE

A variety of quantum algorithms are developed that show quantum speedup in comparison to classical ones (Grover, 1996; Kitaev, 1995; Shor, 1997). Other proposed algorithm perform tasks such as finding eigenvalues and eigenvectors (Abrams and Lloyd, 1999), or solving linear systems of equations (Harrow *et al.*, 2009). Normally, in all these algorithms an input problem must be prepared. This preparation can be computationally costly and be a hidden challenge.

Challenge 1.- Developing quantum algorithms for data compression that represent exponential data with a polynomial circuit depth and polynomial number of qubits.

One first challenge is the quantum state preparation (QSP). One area where quantum computers are anticipated to be impactful and provide computational advantage is in the field of Machine Learning (ML), where classical data are processed, such as data classification. Therefore, a first component of any quantum algorithm with ML purpose inevitably deals with loading classical data into the quantum memory of the qubits (Araujo *et al.*, 2021; Zhang *et al.*, 2022). While an exponential data can be saved on a linear number of qubits, the circuit depth is generally exponential.

Challenge 2.- Software stack that facilitates the representation of classical many-body problems, such as in fusion energy science, in terms of qubit operators.

A second challenge is in the preparation of a given problem in terms of a suitable Hamiltonian. What is normally not discussed in literature is the cost of preparation of the input problem in terms of proper input Hamiltonian, e.g., in terms of a linear combination of Hermitian operators. Some solutions to this challenge already exists. In the field of condensed matter physics or quantum chemistry, computer softwares exist that prepare the input problem (McClean *et al.*, 2019; Qiskit contributors, 2023). In general, there is a lack of tools to create a direct representation of a generic Hamiltonian in the language of quantum computers. A less explored field is simulation of plasma fluid dynamics with quantum computers or quantum-inspired algorithms. Given simulation of plasma fluid dynamics is at the core of industrial business such as General Atomics, a quantum solution can have tremendous impact. Quantum algorithms for solving differential equation exists (Harrow *et al.*, 2009), having a software stack that facilitates the representation of the problem of complex differential equations into spin Hamiltonian is then necessary.

OPPORTUNITY

Priority Research Direction 1.- Developing quantum state preparation algorithm that are inspired by powerful established QIS theoretical techniques (e.g., Tensor Network) to reduce computational resource, discovering quantum advantage.

QIS is supported by powerful classical techniques such as Tensor Network, MERA, and MPS. Nevertheless, direct conversion of the underlying tensors to one- and two-qubit elementary gates comes with exponential cost. As one possible direction of research, some of us recently proposed a different outlook to quantum circuit instruction (Jouzdani *et al.*, 2022). We believe similar considerations can benefit QSP problem in fields such as classical ML.

Priority Research Direction 2.- Tailored software that can be used to prepare generic input problems in terms of qubit operators.

While second quantization is an appropriate way to prepare quantum many-body problem, in most industrial problems particle symmetry is immaterial. The industry could benefit from a software that is used to prepare a larger

[1] Corresponding author, email: jouzdanip@fusion.gat.com

class of problems. A class of problems appears in fusion energy science, where generally the dynamics of a classical many-body system is the target of simualtion.

ASSESSMENT

Comparison with current classical algorithm, and application to real-world problems. The comparison should consider time to solution, including any input preparation, aside from other metrics such as accuracy of the solution and alike.

TIMELINESS OR MATURITY:

Quantum advantage may not necessarily (in short term) be discovered in relation to quantum many-body physics and quantum chemistry problems. The search for quantum advantage should be comprehensive. There are good reasons to believe that some problems in real-world industry can be targeted more efficiently by quantum computers. The first step in this expedition, is efficient tools and algorithms that allows quantum computer to interact with real-world problems and classical data.

REFERENCES

Abrams, D. S., and S. Lloyd (1999), Physical Review Letters **83** (24), 5162.
Araujo, I. F., *et al.* (2021), Scientific Reports **11** (1), 6329.
Grover, L. K. (1996), doi.org/10.48550/ARXIV.QUANT-PH/9605043.
Harrow, A. W., A. Hassidim, and S. Lloyd (2009), Physical Review Letters **103** (15), 150502.
Jouzdani, P., C. W. Johnson, E. R. Mucciolo, and I. Stetcu (2022), Phys. Rev. A **106**, 062435.
Kitaev, A. Y. (1995), arXiv:quant-ph/9511026 [quant-ph].
McClean, J. R., *et al.* (2019), arXiv:1710.07629 [quant-ph].
Qiskit contributors, (2023), “Qiskit: An open-source framework for quantum computing,” .
Shor, P. W. (1997), 10.1137/S0097539795293172.
Zhang, X.-M., T. Li, and X. Yuan (2022), Phys. Rev. Lett. **129**, 230504.

Development of an Experimentally-Inspired and Application-Driven Quantum Internet Stack

Dimitrios Katramatos (BNL, dkat@bnl.gov), Julian Martinez-Rincon (BNL, jmartinez1@bnl.gov), Gabriella Carini (BNL, carini@bnl.gov), Tim Stavenger (PNNL, timothy.stavenger@pnnl.gov), Torsten Hoefler (ETH Zürich, v-torhoefler@microsoft.com), Sebastian Will (Columbia University/BNL, sw3151@columbia.edu), Hong Tang (Yale University, hong.tang@yale.edu), Matthew Lahaye (AFRL, matthew.lahaye@us.af.mil), Laura Wessing (AFRL, laura.wessing.1@us.af.mil), Kathy-Anne Soderberg (AFRL, kathy-anne.soderberg@us.af.mil), Eden Figueroa (Stony Brook University / BNL, eden.figueroa@stonybrook.edu).

Topic: Quantum Networking Models and Applications

1. Challenge: Quantum Internet Stack Concept and Development.

Long-distance distributed quantum processing implementations will require the deployment of extended quantum networks delivering entanglement on-demand to the quantum processing nodes. The primary architectural model governing this envisioned '*quantum internet of things*' is the quantum internet stack [1]. In this model, a quantum physical layer formed by entanglement sources, quantum memories, gates, and routers will deliver entanglement to a quantum data link layer defining endpoints where robust entanglement has to be delivered. The efficient long-distance distribution of entanglement using quantum repeaters is governed by a software-defined on-demand entanglement distribution layer. Once the entanglement is distributed, a quantum teleportation layer will reliably provide on-demand teleportation services among the quantum processing nodes. Finally, teleportation-based applications will be defined, capable of running distributed versions of quantum algorithms across the quantum internet of things. This architecture can be seen on the right column of the figure below.

In this '*quantum internet of things*' model, on-demand quantum operations within a quantum network, such as qubit generation or Bell state measurements, will require the precise time synchronization and remote control of classical devices over ancillary classical network infrastructure. The classical network oversees the flow of control commands responsible for the execution of quantum network operations. Furthermore, the classical network needs to initiate and coordinate the operations of quantum applications that require quantum communication between quantum devices. Therefore, the stack needs to enable such communications allowing applications to programmatically access quantum communication functionality.

2. Opportunity: Create an Experimentally-Inspired and Application-Driven Hybrid Quantum Internet Stack.

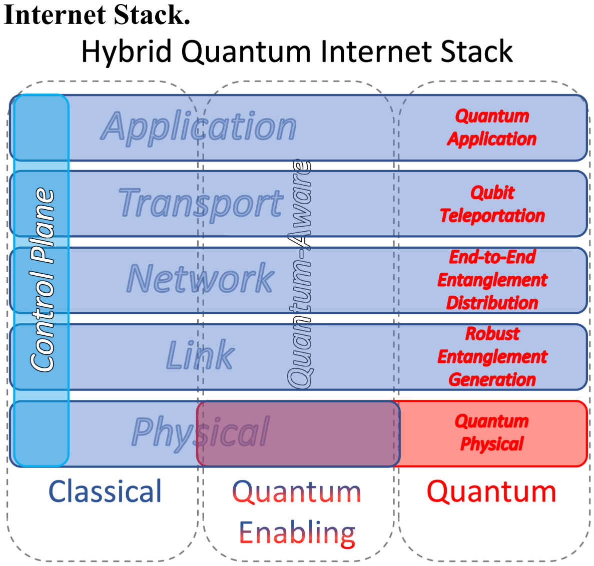

The complexity in orchestrating the tasks required in the quantum internet of things concept defined above calls for a new quantum-aware control paradigm. We envision an evolution of the classical network stack to accommodate quantum-enabling and quantum operations simultaneously. The components and functionality of this hybrid stack can be separated into three stack protocol sets, as depicted in the figure left: (i) the quantum stack, which will be a set of protocols controlling interconnected quantum devices capable of performing the key operations for entanglement generation and distribution, (ii) the quantum-enabling stack, which will be a set of quantum-aware protocols enabling the interconnection of quantum and classical devices in a network specifically built for maintaining quantum coherence (e.g., polarization or time-bin superpositions) carried by photons defined within restrictive temporal envelopes, and (iii) the existing classical stack that will be used to connect with quantum-aware hardware to control quantum network operations. Furthermore, the control plane should also be quantum-aware and comprise a set of dedicated protocols within the classical and the quantum-enabling stacks required to control and orchestrate operations at the device, node, domain, and network level. At the application layer, we envisage the need for standardized APIs for enabling quantum applications to utilize quantum resources over the hybrid network. A set of quantum communication primitives that abstract the quantum communication capabilities and hide their inner workings and complexity will enable the development of high-level programmatic environments that can be utilized in developing quantum applications.

3. Assessment: Emulation and Reference Implementation

To evaluate the feasibility of our architecture, it will be necessary to develop quantum network emulators that are based on the hybrid quantum internet stack model. A comparison of the emulator results against theoretical results will verify the correctness of its operation. Studying the effects of multiple parameters on the operation of the stack-driven quantum networks, such as propagation distance, entangled generation rates, photon loss rates, and quantum decoherence times will allow us to benchmark the operation of multi-node quantum networks. Another avenue to evaluate the architecture model will be to use the quantum internet stack to extend existing classical distributed computing communications standards such as the widely used Message Passing Interface (MPI) [2] into a quantum version (QMPI), offering an API that can utilize both classical and quantum resources at the behest of applications [3]. For first realizations, the implementation details of QMPI are likely to be unique to each hardware system, however, part of the assessment will be to eventually create a hardware-agnostic standard. A similar approach could be used with another common parallel programming model, the Partitioned Global Address Space (PGAS), which acts as remote, shared memory access with local affinity. A Quantum PGAS (QPGAS) would be quantum-teleportation-aware and have a strong correlation to distributed quantum algorithms that are typically designed around manipulating shared physical quantum resources over message communications. Using a quantum network emulator as a platform, it will be straightforward to develop reference implementations of communication standards, such as QMPI or QPGAS on top of the hybrid stack, to test, validate, verify, and demonstrate the completeness of the proposed APIs. If these APIs are shown to be complete, new standards could be proposed for broader acceptance.

4. Timeliness or Maturity: Hybrid Quantum Internet Stack Realization

First iterations of this hybrid stack concept should be experimentally tested using controllable quantum nodes in quantum internet testbeds across the United States [4-6]. For example, the BNL/Stony Brook quantum internet testbed or the Illinois Express quantum networking testbed are envisioned to be controlled over the classical network using first instances of a Quantum-Aware Control Plane (QACP). Such QACPs are envisioned to orchestrate the operations of various devices over the network infrastructure to enable quantum information flow and will be built on a series of drivers and software that will control system elements over the network, acting as a unified control language. We envision such implementations to achieve maturity in the next five years. Looking ahead into the next decade, similar approaches could be followed to control the entanglement distribution and quantum teleportation services envisioned in the general quantum internet stack. With this at hand, first instances of high-level implementations of QMPI or QPGAS and related applications could be realized within the next decade.

5. References

1. NPJ Quantum Information 8, 121, 2022.
2. MPI: a standard message passing interface. Supercomputer 12, 56–68 (1996).
3. Distributed quantum computing with QMPI. Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. SC ’21. Association for Computing Machinery, doi: 10.1145/3458817.3476172, (2021).
4. D. Du and et al., An elementary 158 km long quantum network connecting room temperature quantum memories, arXiv:2101.12742, (2021).
5. Reconfigurable Quantum Local Area Network Over Deployed Fiber, PRX Quantum **2**, 040304 (2021).
6. Illinois Express Quantum Network (IEQNET): Metropolitan-scale experimental quantum networking over deployed optical fiber, arXiv:2104.04629v1 (2021).

- **Title:** Integrating devices, protocols, simulation and experiments research to co-design a scalable quantum network.
- **Authors (with affiliations and email addresses):** Rajkumar Kettimuthu, Xu Han and Joaquin Chung, Argonne National Laboratory, {kettimut, xu.han, chungmiranda}@anl.gov
- **Topic:** Models, codesign and integration
- **Challenge:** Though proof-of-concept quantum communication over optical fibers has been demonstrated at lab, campus and metropolitan scales, a number of open questions and challenges remain in quantum networking. As experimentalists develop quantum memories, fast entanglement sources, low-noise detectors, frequency converters, and repeaters, how do we evaluate them at scale in terms of number and heterogeneity of devices, distance between devices, volume of information transmitted, the number of users and the diversity of applications? What abstractions and interfaces do different kinds of devices need to provide? How do we evaluate various network architectures (for example, centralized versus decentralized control, in-band versus out-of-band signaling) and protocol stacks at scale? Some quantum network testbeds exist and there are efforts to build new testbeds but they have several limitations. For example, these testbeds are campus or metropolitan scale and they have only limited number of devices as many key devices are still in research stage. Simulations can provide better scalability and much more flexibility while being cost-effective, but accurate models and appropriate level of abstraction are critical for achieving the desired fidelity. There are things that can be learnt from the classical network simulation which has a rich history ranging from packet-level discrete event simulators to approaches that employ a coarse resolution than packet-level simulation to hybrid approaches. At the same time, quantum networks are very different from their classical counterparts. Two or more quantum states can become entangled, and local operations can have nonlocal effects. For example, multipartite entanglement allows creation of routing protocols that do not have classical analogs. In fact, there has been efforts underway to build quantum network simulators but an integrated effort involving researchers building devices, theorists developing protocols, system software people and experimentalists is required to move the state-of-the-art in quantum networking.
- **Opportunity:** Progress is being in quantum networking devices and components [1,2,3] such as quantum repeaters, quantum memories, quantum frequency converters, entangled photon sources and single photon detectors. Likewise, several quantum network architectures and protocols are being developed [4,5,6]. Testbeds at campus and metropolitan scale are being actively built [7,8,9]. Several Quantum network simulators are available now and are being actively enhanced [10,11,12]. Quantum applications remain in the realm of theory, usually at circuit level description and thought experiments. To achieve scalable quantum communications, we require a systems approach that brings together devices, protocols, simulation and experiments to codesign the quantum networking stack. With active research (although in a siloed fashion) in all the individual areas, there is a great opportunity now for an integrated research program to realize a practical quantum network.
- **Assessment:** Success involve making big strides in all the layers in a quantum networking stack ranging from devices to applications with the protocols and network architectures in between. Significant progress in developing abstractions and interfaces is also important. One potential metric is the advancement in the scale of quantum networks along different dimensions such as distance between nodes, heterogeneity of devices connected, number of devices connected.
- **Timeliness or maturity:** The fact that a) several quantum networking demonstrations has been done at the campus and metropolitan scales recently; b)numerous research activities are ongoing in the devices front; c) a handful of testbeds are being built; and d) simulation efforts have produced some reasonable results, makes it timely to ramp up a co-design and integrated approach to advance the quantum networking stack is timely now. The impact of success could be big resulting in moving quantum network from metro scale to regional and national scale.

- **References**

[1] F. Rozpedek et al. Quantum repeaters based on concatenated bosonic and discrete-variable quantum codes. arXiv preprint arXiv:2011.15076, 2020.
[2] A. Dibos et al. Purcell Enhancement of Erbium Ions in TiO2 on Silicon Nanocavities, Nano Letters 2022 22 (16), 6530-6536 DOI: 10.1021/acs.nanolett.2c01561
[3] Han, X., Fu, W., Zhong, C. et al. Cavity piezo-mechanics for superconducting-nanophotonic quantum interface. Nat Commun 11, 3237 (2020). https://doi.org/10.1038/s41467-020-17053-3
[4] J. Chung et al. "Design and Implementation of the Illinois Express Quantum Metropolitan Area Network," in IEEE Transactions on Quantum Engineering, vol. 3, pp. 1-20, 2022, Art no. 4100920, doi: 10.1109/TQE.2022.3221029.
[5] K. Kapoor et al. "Picosecond Synchronization System for the Distribution of Photon Pairs through a Fiber Link between Fermilab and Argonne National Laboratories," in IEEE Journal of Quantum Electronics, doi: 10.1109/JQE.2023.3240756.
[6] A. Zang et al. Entanglement Distribution in Quantum Repeater with Purification and Optimized Buffer Time, IEEE Infocom Workshop on Network Science for Quantum Communication Networks, May 2023.
[7] New quantum loop provides long national testbed for quantum communication technol- ogy. https://www.anl.gov/article/new-quantum-loop-provides-long-national-testbed-for-quantum-communication-technology, 2020.
[8] J. Chung et al. Illinois Express Quantum Network (IEQNET): Metropolitan-scale experimental quantum networking over deployed optical fiber. arXiv preprint arXiv:2104.04629, 2021.
[9] D. Du et al. An elementary 158 km long quantum network connecting room temperature quantum memories. arXiv preprint arXiv:2101.12742, 2021.
[10] X. Wu et al. SeQUeNCe: a customizable discrete-event simulator of quantum networks, 2021 Quantum Sci. Technol. 6 045027 DOI 10.1088/2058-9565/ac22f6
[11] Matsuo, T. Simulation of a dynamic, RuleSet-based quantum network. arXiv:1908.10758 (2020).
[12] Coopmans, T., Knegjens, R., Dahlberg, A. et al. NetSquid, a NETwork Simulator for QUantum Information using Discrete events. Commun Phys 4, 164 (2021). https://doi.org/10.1038/s42005-021-00647-8

- **Title:** Scalable interconversion of electronic and photonic quantum bits for practical quantum computing and communication systems
- **Authors**: Pedram Khalili, Department of Electrical and Computer Engineering, Northwestern University, pedram@northwestern.edu
  Prem Kumar, Departments of Electrical and Computer Engineering and Physics, Northwestern University, kumarp@northwestern.edu
- **Topic:** Codesign and Integration
- **Challenge:** While most practical qubit operations are based on electrical gates and controls, coherent transmission of such qubits over long distances is not practical. On the other hand, optical photons traveling through fibers or open space can maintain their quantum states across long distances. Unfortunately, conversion of qubit signals at microwave bands to optical-frequency photons with high efficiency and fidelity has been a major challenge due to the large wavelength mismatch between the two. Solving this problem is essential for the development of practical quantum computing and communication systems where distant qubits can communicate coherently via optical photons that preserve their quantum states. Addressing this challenge requires the development of robust, efficient, scalable, and integrated microwave to optical photon transducers. While progress has been made in recent years in the development of transducers based on hybrid platforms involving optical waveguides and magnonic, phononic, and electro-optical devices (*1-11*), most of this work has not ventured beyond fundamental physics or single-device-level studies. Developing practical transducers with industry-relevant scale and complexity that provide quantum links between qubits across long distances remains an unsolved challenge. Addressing this challenge will require a codesign approach across the technology stack from materials, device engineering, integration, and system design (Figure 1).
- **Opportunity:** The increasing progress and availability of capabilities for design, synthesis, and heterogenous integration of magnonic and phononic materials (e.g., ferro, ferri, and antiferromagnets supporting magnons, and ferroelectric materials supporting electrically tunable phonon modes) creates a timely opportunity to address this challenge. As an example, magnons—quanta of spin waves in a magnetically ordered material—are a possible contender as an intermediary particle of quantum frequency converters. Their frequency, which can be designed by controlling the exchange, magneto-crystalline anisotropy, and shape of the magnet, is typically in the ~1-30 GHz range. On the other hand, their wavelengths are much shorter than the free-space wavelength of the same frequency (~100 nm – 10 μm) and can be engineered by designing films or multilayers with specific dispersion characteristics. This allows for small devices with optimized coupling to both microwave and optical photons. In addition, recently discovered interfacial magneto-electric and spin-orbitronic effects create new opportunities to develop devices with high conversion efficiencies between microwave photons and the magnons (*12-15*). Finally, a wide range of ferromagnetic and antiferromagnetic materials are now routinely integrated into existing semiconductor manufacturing processes of leading foundries, largely due to their applications in (classical) memory and sensing products (*16*). Examples are TSMC, Samsung, and GlobalFoundries 22 nm and 28 nm nodes which all have embedded magnetic memory (MRAM) options, with sub-20 nm nodes in development. Hence, their synthesis, processing, and large-scale integration build on a large pool of knowledge and prior investments by industry, government, and academia. Integration of quantum transducers within advanced foundry nodes will enable them to be placed in tight proximity to advanced classical computing, control, and interface circuitry needed for large-scale quantum computation and communication. We expect that this will allow a great leap forward in terms of the complexity and industry/application-relevance of the circuits being prototyped. This will be comparable to the effect that the increased availability of silicon photonics in semiconductor foundries has already had on quantum technology development in recent years (*17*).

- **Assessment:** The outcome of this research effort should be a robust, energy-efficient, and integrated microwave to optical photon transducer. The key metrics for success will be high conversion efficiency and fidelity, integration on a monolithic or chiplet platform to achieve high qubit density, and qubit conversion throughput.

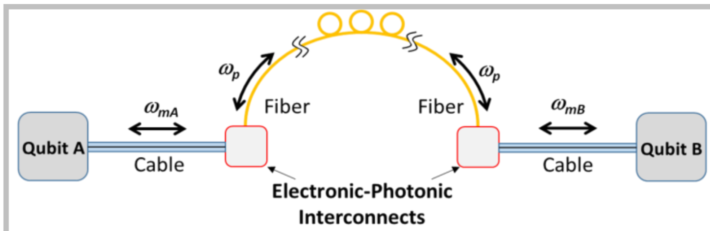

Fig. 1. Quantum interconnects that provide efficient, coherent, and scalable interconversion of electronic and photonic quantum signals can be used to couple multiple electrically accessible qubits via long-distance fibers. This is an essential ingredient for large-scale quantum computing and communication systems.

- **Timeliness or maturity:** This is a timely topic due to two reasons: 1. The significant progress over the past decade in developing microwave-frequency quantum computing platforms. Further scaling of these technologies requires true quantum coupling of separate quantum computers to address larger computational problems; and 2. The tremendous progress in integration of various ferroic materials within established semiconductor manufacturing, which provides a pathway to scalable integration of large transducer arrays in commercial foundries.
- **References**

1. Y. Tabuchi *et al.*, Hybridizing Ferromagnetic Magnons and Microwave Photons in the Quantum Limit. *Physical Review Letters* **113**, 083603 (2014).
2. X. Zhang, N. Zhu, C.-L. Zou, H. X. Tang, Optomagnonic Whispering Gallery Microresonators. *Physical Review Letters* **117**, 123605 (2016).
3. R. W. Andrews *et al.*, Bidirectional and efficient conversion between microwave and optical light. *Nature Physics* **10**, 321-326 (2014).
4. J. Bochmann, A. Vainsencher, D. D. Awschalom, A. N. Cleland, Nanomechanical coupling between microwave and optical photons. *Nature Physics* **9**, 712-716 (2013).
5. A. Vainsencher, K. J. Satzinger, G. A. Peairs, A. N. Cleland, Bi-directional conversion between microwave and optical frequencies in a piezoelectric optomechanical device. *Applied Physics Letters* **109**, 033107 (2016).
6. L. Fan *et al.*, Superconducting cavity electro-optics: A platform for coherent photon conversion between superconducting and photonic circuits. *Science Advances* **4**, eaar4994 (2018).
7. B. Bhoi, B. Kim, J. Kim, Y.-J. Cho, S.-K. Kim, Robust magnon-photon coupling in a planar-geometry hybrid of inverted split-ring resonator and YIG film. *Scientific Reports* **7**, 11930 (2017).
8. R. Hisatomi *et al.*, Bidirectional conversion between microwave and light via ferromagnetic magnons. *Physical Review B* **93**, 174427 (2016).
9. X. Zhang, N. Zhu, C.-L. Zou, H. X. Tang, Erratum: Optomagnonic Whispering Gallery Microresonators [Phys. Rev. Lett. 117, 123605 (2016)]. *Physical Review Letters* **121**, 199901 (2018).
10. N. Zhu *et al.*, Waveguide cavity optomagnonics for microwave-to-optics conversion. *Optica* **7**, 1291-1297 (2020).
11. D. D. Awschalom *et al.*, Quantum Engineering With Hybrid Magnonics Systems and Materials. *IEEE Transactions on Quantum Engineering*, 1-1 (2021).
12. J. Zhu *et al.*, Voltage-Induced Ferromagnetic Resonance in Magnetic Tunnel Junctions. *Physical Review Letters* **108**, 197203 (2012).
13. R. Khymyn, I. Lisenkov, V. Tiberkevich, B. A. Ivanov, A. Slavin, Antiferromagnetic THz-frequency Josephson-like Oscillator Driven by Spin Current. *Scientific Reports* **7**, 43705 (2017).
14. A. Slavin, V. Tiberkevich, Nonlinear Auto-Oscillator Theory of Microwave Generation by Spin-Polarized Current. *IEEE Transactions on Magnetics* **45**, 1875-1918 (2009).
15. S. Cherepov *et al.*, Electric-field-induced spin wave generation using multiferroic magnetoelectric cells. *Applied Physics Letters* **104**, 082403 (2014).
16. W. J. Gallagher *et al.*, in *2019 IEEE International Electron Devices Meeting (IEDM)*. (2019), pp. 2.7.1-2.7.4.
17. See, for example, https://psiquantum.com/; https://www.xanadu.ai/photonics.

Quantum Programming 2.0: Hybrid classical-quantum programming in heterogeneous execution environments

David Ittah, Josh Izaac, Lee J. O'Riordan, Nathan Killoran*
Xanadu Quantum Technologies

April 28, 2023

Topics: Computing and programming models, compilation, codesign and integration

Challenge

The past decade has witnessed quantum computing hardware becoming widely publicly available, leading to an emergent "Quantum Programming 1.0" software stack. With small variations, the standard is dominated by a few commonalities: Python-based libraries for users to craft their algorithms, built-in quantum simulators (created in compiled languages like C++), online APIs which accept text-based representations of circuits, tight constraints on what a "quantum program" is, and substantial classical computation often executing in a user's local environment.

This model has enabled things that have never been possible before, but we should not expect that we have gotten the entire quantum programming stack right in a single technological generation. These present-day patterns may prove insufficient for new sets of demands, which are already starting to appear. Freshly emerging challenges and limitations will force us to rethink the foundations of the first-generation quantum software stack.

One challenge is supporting the growing diversity in what "quantum programming" means. Many established and emerging algorithms have, by necessity, some hybrid nature to them: error correction, error mitigation, circuit cutting, quantum gradient computation, variational algorithms, etc. Arguably, even textbook algorithms like Shor's are better captured with mixed quantum and classical instructions (e.g., classical control flow). Drawing arbitrary programmatic restrictions around quantum and classical operations will make quantum programs less flexible, less modular, and less hackable, slowing the pace of research overall. Similarly, there are a number of devices and accelerators available where a user may want to run different parts of their hybrid program. How do we allow hybrid computations to be intelligently distributed amongst QPUs, CPUs, and GPUs?

As a final challenge, while the accessibility and scientific ecosystem of Python are important for researchers, Python is a notoriously slow programming language for the actual execution of complex programs. Dedicated quantum simulators written in compiled languages only partially mitigate this, since they focus exclusively on quantum operations—classical parts of a user's program still execute in Python. How can we leverage the researcher-friendliness of interpreted languages like Python to capture a user's program, while smoothly moving all execution to environments which are better suited for both classical and quantum instructions?

Opportunity

There exists a major opportunity to reimagine a new "Quantum programming 2.0" stack end-to-end. This new paradigm could be distinguished from its predecessor by putting solutions to the above mentioned

*nathan@xanadu.ai

challenges at its heart: full-stack treatment of hybrid programs, flexible heterogeneous program execution, and staging the execution of entire classical-quantum programs to occur outside of Python. This does not need to be a complete tear-down and rebuilding of our existing stack, but an evolutionary process that keeps the best patterns while substantially improving the weak points.

Upgrading the quantum programming model would deliver substantial benefits to quantum researchers, accelerating the pace of development. Example opportunities include: faster program execution via just-in-time hybrid compilation, joint classical-quantum transforms/passes for program optimization, more seamless integration with modern classical software tools, and removing network communication bottlenecks from quantum programming workflows.

Seizing this opportunity will require the exploration of new ideas and the development of new tools and technologies. It will necessitate broader collaboration across industry, e.g., developing these next-gen standards, and codesign of quantum hardware APIs and quantum software clients. It will also need greater collaboration between industry and academia, e.g., understanding and removing researchers' biggest barriers, and developing additional ideas and algorithms for hybrid programs.

Assessment

Any software stack should be assessed by what benefits it brings the user. With this in mind, successful development of a Quantum Programming 2.0 stack could take the following form: By 2025, a researcher with no special programming expertise can develop a new idea for a quantum algorithm in a Python-based quantum programming library within minutes. The user's entire program (all classical and quantum instructions) is captured from Python, passes through a compiler pipeline that jointly optimizes for both quantum and classical efficiency, and is staged for execution in a high-performance runtime environment. Execution of the user's program is heterogeneous, taking place concurrently across multiple devices including CPUs, GPUs, and QPUs, with minimal overhead nor need for special expertise from the user. This entire stack is modular, extensible, and hackable, allowing third-party developers to build their own applications on top seamlessly. Quantum hardware can register its own specific compilation passes for this stack dynamically, allowing newly available devices to easily integrate.

Timeliness or maturity

A first-gen quantum programming model has already coalesced. Thanks to this, we now have visibility into which parts of this model are (or will soon be) insufficient. Embryonic next-gen ideas are now appearing, such as in the Quantum Intermediate Representation [1], Xanadu's Catalyst [2], Nvidia's CUDA Quantum [3], or the IBM Qiskit Runtime [4]. But these budding ideas will need dedicated time and effort to develop into a clear technological trend. While certain trends may appear organically, there is a risk that a restrictive or uncoordinated development could delay or prevent key ideas from emerging in a timely manner (or at all), slowing the development of quantum technologies overall. As a best-case scenario, a concerted effort to develop the next-gen quantum software stack unlocks substantial capabilities for quantum researchers, leading to a large acceleration in the timeline where useful quantum computers are available.

References

[1] *QIR Specification.* `https://qir-alliance.org`.

[2] *Catalyst.* `https://github.com/PennyLaneAI/catalyst`.

[3] *CUDA Quantum.* `https://github.com/NVIDIA/cuda-quantum`.

[4] *Qiskit Runtime* `https://github.com/Qiskit/qiskit-ibm-runtime`.

Title: Quantum Modeling and Optimization Design Environment (Q-MODE)

Authors:
Yousu Chen, Pacific Northwest National Laboratory, yousu.chen@pnnl.gov
Amin Khodaei, University of Denver, amin.khodaei@du.edu
Rozhin Eskandarpour, Resilient Entanglement, rozhin.eskandarpour@resilastententanglement.us

Topic: Design and development of an application environment for quantum model and algorithm integration

Challenge: Quantum computing has the potential to revolutionize many areas of science and technology, from cryptography and drug discovery to optimization and machine learning. As the field matures, the need for sophisticated tools that can model and simulate quantum computing problems becomes more critical. Despite the progress made by existing quantum computing frameworks, there is still a need for a comprehensive, user-friendly environment for the design, optimization, and integration of quantum systems. Such modeling systems exist for optimization problems in the classical domain, e.g., GAMS and AIMMS. However, this is still lacking in the quantum domain. The challenges facing the development of a quantum computing environment like Q-MODE include:

- Creating an expressive, user-friendly domain-specific language that allows researchers and practitioners to quickly and easily model optimization problems using quantum computing.
- Ensuring compatibility with existing quantum computing frameworks, libraries, and hardware.
- Developing a cohesive modeling environment that integrates quantum and classical components, enabling seamless codesign and optimization of hybrid systems.

Opportunity: Q-MODE can address these challenges by offering a comprehensive quantum modeling and optimization design environment tailored to the unique requirements of quantum computing. The key features and benefits of Q-MODE include:

- A new, expressive domain-specific language specifically designed for quantum computing, enabling users to easily describe quantum circuits, gates, algorithms, and select the desired optimization algorithm.
- Integration with popular quantum computing frameworks such as Qiskit, Cirq, and PennyLane, allowing users to leverage the latest quantum computing resources and advancements.
- Support for a wide range of optimization techniques and solvers, including quantum annealing, QAOA, VQE, and hybrid classical-quantum optimization methods.
- A user-friendly interface and visualization tool for designing, simulating, and optimizing quantum systems, making it accessible to researchers and practitioners with varying levels of expertise in quantum computing.

- Comprehensive documentation, tutorials, and examples to help users get started with Q-MODE and explore its full capabilities.
- Active community engagement and support, promoting collaboration and knowledge-sharing among users and developers.

Assessment: The success of Q-MODE will be evaluated based on the following criteria:

- Adoption and usage: Q-MODE should be embraced by the research community and industry partners, as evidenced by the number of users, publications, collaborations, and real-world applications utilizing the software.
- Compatibility and integration: Q-MODE must maintain compatibility with existing quantum computing frameworks, libraries, and hardware, enabling users to leverage the latest resources and advancements in the field.
- Ease of use: The domain-specific language, interface, and tools provided by Q-MODE should be user-friendly, accessible, and efficient, making it easy for researchers and practitioners to model, simulate, and optimize quantum systems.

Timeliness or Maturity: Quantum computing is at a critical juncture, as advances in hardware and software have brought us closer to practical applications and breakthroughs. The maturation of the field necessitates a robust, user-friendly modeling and optimization environment that can support the growing needs of the research community and industry partners. As a result, the development of Q-MODE is both timely and vital for the quantum computing ecosystem. By enabling a widespread application of quantum computing, Q-MODE can provide a powerful tool for the next generation of quantum researchers and practitioners, enabling them to tackle more complex and ambitious problems. The success of Q-MODE will not only streamline and enhance the quantum computing workflow but also contribute to the broader adoption of quantum computing technologies across various industries, accelerating innovation and potentially leading to transformative breakthroughs.

References:

[1] IBM Qiskit, "Qiskit: An Open-source Framework for Quantum Computing," https://qiskit.org/.
[2] Google AI Quantum, "Cirq: A Python Library for Quantum Circuits," https://github.com/quantumlib/Cirq.
[3] Xanadu, "PennyLane: Quantum Machine Learning, Automatic Differentiation, and Optimization," https://pennylane.ai/.
[4] J. Preskill, "Quantum Computing in the NISQ Era and Beyond," Quantum, vol. 2, p. 79, 2018.

Geometric Dual of Non-Unitary Qubit Devices

Lucas Kocia (lkocia@sandia.gov), Shivesh Pathak
Sandia National Laboratories

Topic: Quantum computing applications, models, algorithms development and error correction and mitigation.

Challenge: Since quantum devices are generally susceptible to non-unitary noise in addition to their intended unitary dynamics, their qubits evolve under non-unitary maps and must generally be error-corrected to become useful for practical applications. Our current paradigm of characterizing the usefulness of qubits based on error-correction requires them to be on the surface of their Bloch sphere and non-unitary noise contracts the Bloch sphere.

There are a lot of outstanding challenges in understanding, simulating, and characterizing non-unitary qubit dynamics. For instance, Lindbladian models are often unwieldy, difficult to understand, and generally unscalable, while benchmarking characterizations of quantum devices are difficult to translate to a measure of the device's algorithmic usefulness. At the application-level, an enormous amount of overhead must also be expended to implement quantum error correction. At the algorithmic level, quantum algorithms that use non-pure mixed states, such as those related to Gibbs sampling, are among the most poorly understood and unscalable. This is also true of quantum algorithms for high-energy physics, quantum field theories, and general relativity.

Opportunity: The relationship between general relativity and quantum theory was greatly extended over twenty years ago by the AdS/CFT correspondence conjecture between anti-de Sitter (AdS) spacetime and conformal field theories (CFTs). This relationship establishes a duality between generally non-unitary qubit dynamics and geometric or gravity geodesic equations of motion, among many others that form a sort of "dictionary". This correspondence manifests the holographic principle: the properties of a larger dimensional bulk geometry are encoded in a lower-dimensional field theory. Recently, these correspondences have produced particularly promising results regarding the thermodynamic and entropic properties of black holes on the gravity side, and quantum error correction on the quantum side. Surprisingly, the two phenomena are often found to be dual to each other.

Using the algebraic formalism provided by CFTs, we can see that when qubit-based devices cannot offer quantum advantage, non-unitary qubit operations can only be written in terms of finite-dimensional Witt algebra generators that produce a cover of the conformal symmetry group $SO^+(2,2)/SO^+(1,2)$ under exponentiation. When they can offer quantum advantage, this algebra possesses a non-zero "central" charge and can be extended to an infinite-dimensional Virasoro algebra. This allows any (hypothetical or actual) quantum device to be expressed in terms of the conformal symmetry operations that it can produce and to determine how much this central extension is used.

However, there appear to be no algorithms that explicitly make use of the non-compact and non-unitary subspace of the full $SO^+(2,2)/SO^+(1,2)$ group. This is largely because it is not possible to define a finite logical qubit basis on a non-compact space and so only the compact $SU(2^N)$ subspace is used. While all practical quantum devices technically explore the non-unitary part because they experience noise, this is viewed as something to be mitigated instead of leveraged; qubits must remain in $SU(2^N)$ to remain useful for current algorithmic implementations.

There is a potentially significant missed benefit to not making use of the full space. $SO^+(2,2)/SO^+(1,2)$ allows two-dimensional qubits to be reformulated in a larger-dimensional dual space. This relates their angular rotational, dilatation and other degrees of freedom to the shifts and boosts of curved spaces that can be solved by classical general relativity when the central charge is zero (i.e. no quantum speedup). With only the unitary part of $SO^+(2,2)/SO^+(1,2)$, this correspondence only generates comparatively simpler Euclidean geometries and not the full diversity of Riemannian geometries. In this sense, current unitary approaches to quantum algorithm development cannot efficiently simulate many problems in general relativity and this may be why many current quantum algorithms for high-energy scale particularly badly.

A finite logical qubit basis can be defined by *compactifying* the non-compact subspace of $SO^+(2,2)/SO^+(1,2)$ onto a compact space, such as another qubit. Furthermore, the AdS/CFT correspondence has developed approximate error-correcting codes that correspond to operator algebraic prescriptions for erasure-protecting codes. It has also been shown that depolarization errors can frequently be converted to erasure errors which subsequently improves the threshold distance of the quantum error-correcting code. Geometric dual notions such as Bogoliubov transformations, location in the radial direction, and the holographic entropy bound reexpress notions from quantum error correction. Black hole scrambling dynamics that conceals quantum information by encoding it non-locally has been shown to also protect it as a byproduct like a quantum error-correcting erasure code. Studying quantum error-correction in this dual geometric/gravity formulation may allow us to better leverage the non-unitary potential intrinsic in today's devices.

Assessment: Determining how many units of the full $SO^+(2,2)/SO^+(1,2)$ group are necessary in quantum algorithms and can be generated by quantum devices will more accurately align with their useful benefit over classical devices and algorithms and generically lower the amount of qubits necessary. As a result, it would be extremely useful to develop a two-dimensional conformal "volume" metric, similar in principle to the quantum "volume" used by IBM, but directly measuring a device's utility to outperform classical computers *at the algorithmic level* and therefore advance the frontiers of computational science.

Potential solutions to preserving the non-compact subspace of qubits would likely initially involve extensions of quantum erasure codes or extensions of standard quantum correction codes to some approximate basis over the larger non-compact subspace, or a compactification of one qubit onto more qubits. Success could be evaluated with respect to the above-mentioned conformal volume metric. Moreover, the measure of success at the algorithmic level would follow a similar route; expressing our current resource estimates in terms conformal units, instead of unitary T gates or similar, would likely be especially beneficial for algorithms of high-energy physics.

Timeliness or maturity:

Many of the links between the geometric/gravity dual formulation of non-unitary qubit dynamics and quantum error correction have only recently been made. Moreover, despite the invention of quantum error correction following the AdS/CFT correspondence conjecture, the two communities have only begun to be bridged. At the application level, now is also in some sense the perfect time since intermediate-scale near-error correcting devices will likely approaching the sweet spot of conformal control over their non-purity, i.e. approximate error correction will in a sense be inherently the best that we will be able to do with these devices.

Furthermore, the dual reformulation of non-unitary qubits in terms of curved spaces has the potential of helping us develop fundamental physical limits on quantum processors by allowing for speed and scrambling limits from classical general relativity to be applied to the quantum setting. The tools developed in that community have never been used to benchmark quantum device capabalities and we suspect that they are particularly capable of doing so at the algorithmic level for near-term devices that are naturally non-unitary.

Scaling up NISQ era simulations through robust hybrid computing

Nicholas P. Bauman[(a)], Eric J. Bylaska[(a)], Daniel Claudino[(b)], Niladri Gomes[(c)], Karol Kowalski[(a)] *, Ang Li[(a)], Bo Peng[(a)]

[(a)]Pacific Northwest National Laboratory, [(b)]Oak Ridge National Laboratory,

[(c)]Lawrence Berkeley National Laboratory

* karol.kowalski@pnnl.gov

Topic: Defensible demonstrations of quantum computing as a means to simulating nature are contingent not only on the advances in the design of quantum hardware but also on the ability to adopt new methodological advances in the theory of correlated many-body systems and elements of high-performance computing (HPC). Among the desired features of these methodologies, one should list the possibility of utilizing hybrid computational resources integrating the most appealing features of quantum and classical computing and the ability to construct flexible workflows that can naturally adapt to ever-evolving quantum computing technologies and provide much-needed tools for transitioning from the noisy intermediate-scale quantum (NISQ) era to fully fledged error-corrected quantum computing. Close integration between existing exascale platforms and quantum hardware has the potential to lead to more realistic simulations of chemical processes. The "center of mass" in these simulations can be shifted towards quantum computers as quantum computing matures. Therefore, our position paper focuses on components relevant to this strategy: applications, models, algorithms, and codesign/integration aspects.

Challenge: The current applications of quantum computing dominated by NISQ devices are limited by severe bottlenecks that, if not adequately addressed, may lead to a stalemate situation in enabling realistic quantum simulations for chemical systems. Among the most pressing issues, one should mention: (1) circuits that are too-complicated/deep/over-parameterized for a straightforward encoding of physically meaningful information to describe states of molecular and periodic systems in various basis sets (Gaussian vs. plane waves), (2) the lack of quantum algorithms to effectively reduce the quantum problem's dimensionality through the spatial, temporal, and energy scales or appropriate discretization techniques, which leads to the excessive number of parameters that need to be used to attain chemical accuracy, and (3) the number of measurements needed grows excessively large, defining in a natural way the system-size limits tractable by quantum simulations. These three challenges, in addition to inherent problems with the NISQ devices (associated with inefficient/inexistent error-correcting procedures), pose a significant challenge for further advances in quantum computing for chemical systems.

Opportunity: To bypass these problems, we advocate for hybrid computing, where classical computing is integrated with the efficient form of distributed or parallel quantum computing. Thirty years ago, a similar strategy for classical computing (distributed memory models) paved the way for fully-fledged exascale simulations. The currently existing exascale architectures can perform many-body simulations with 10^{11}-10^{12} parameters, a fact that cannot be ignored in designing more efficient and accurate hybrid computing workflows that utilize quantum computing kernels. We envision that such workflows, in a long-time perspective, will provide mechanisms to adapt to evolving quantum technology and will lead to well-defined areas of applications that require pre-and post-processing, usually related to the preparatory steps for calculating effective representation of the quantum problems, optimization of the measurement process, and implementing solvers of non-linear optimization problems.

Although almost all current applications capitalize on some form of hybrid computing, new algorithms that can fully utilize exascale computers and quantum registers composed of hundreds of qubits are needed. An emerging opportunity is associated with (1) a class of methodologies that can be translated into constant-depth algorithms composed of coupled small-dimensionality problems and (2) measurement protocols utilizing various forms of tapering, commutation relations, and unitarizations, which can be executed in parallel using distributed quantum computing.

Co-design processes should streamline hybrid computing by providing mechanism to leverage the existing HPC and quantum computing infrastructure, including HPC tools such as TAMM (Tensor Algebra for Many-body Methods), quantum computing software (for example, XACC and QDK), and mature computational chemistry infrastructure (NWChem, NWChemEx, SPEC). We envision a rapid progress in several resource-efficient quantum algorithms based on: (1) effective Hamiltonian theories, quantum flow equations, and basis set discretization techniques for dimensionality reduction and explicit utilization of interactions locality, (2) quantum Krylov methods, and (3) stochastic approaches based on various representations of Fermionic problem.

Assessment: The success of the hybrid computing strategy should be measured at various levels by reflecting on the potential of quantum/classical simulations. Although achieving success may still be contingent upon unforeseen factors, it is reasonable to expect that the success metric should include the following elements:

- Demonstrate hardware simulations for chemical/physical systems taking advantage of 200-400 qubits and currently existing leadership computing facility architectures.
- Ensure new advances are entirely (or in large part) interoperable with existing infrastructure for computational chemistry and material sciences (quantum chemistry, statistical physics, and periodic systems) and novel HPC simulators.
- Demonstrate the advantages of quantum computing in solving complex processes by addressing problems outside of the reach of conventional computing, including challenging problems involving multiple states in various energy regimes and extensions to periodic systems.

Timeliness or maturity: Based on the current projected developments of quantum hardware, quantum algorithms, and HPC, the discussed strategy elements will enable to by-pass inherent biases of the approximate methods used in chemistry/physics applications and unlock new classes of information in hypothesis-driven quantum simulations. The primary beneficiary of the outlined effort will be quantum chemistry, where quantum computing will result in unprecedented capabilities in modeling processes characterized by strong correlation effects (complex chemical reactions, excited states, and optical properties). However, the universal character of the proposed framework makes it easily transferable to material sciences and nuclear physics.

References

Jérôme F. Gonthier, Maxwell D. Radin, Corneliu Buda, Eric J. Doskocil, Clena M. Abuan, and Jhonathan Romero, Measurements as a roadblock to near-term practical quantum advantage in chemistry: Resource analysis, Phys. Rev. Research 4, 033154

Digital Twin Simulators and Electronic Design Automation for Quantum Hardware: Towards Interoperable and Automated Blueprint Models

Dirk Englund, Massachusetts Institute of Technology, englund@mit.edu
Don Towsley, University of Massachusetts Amherst, towsley@cs.umass.edu
Saikat Guha, University of Arizona, saikat@arizona.edu
Stefan Krastanov[1], University of Massachusetts Amherst, skrastanov@umass.edu

Topic: models, compilation, error correction and mitigation, codesign, integration

Challenge: The development of large-scale quantum computers is hindered by the need for efficient error correction and the limitations of existing hardware. Efficient implementation of quantum algorithms requires (1) effective gate synthesis techniques as well as (2) the integration of quantum technologies across multiple layers of the technology stack and (3) their subsequent optimization. Such holistic design and analysis of quantum systems is further complicated by the lack of (4) standardized physical-layer models and the labor-intensive process of converting research findings into actionable designs.

The unifying theme in these challenges is the lack of good codesign tools for quantum hardware. Classical engineering has faced the challenge of codesign for decades, and has had good success using completely decoupled tools to model various layers of the stack: materials modeling informing lump-element simulations informing simulators of the higher layer of the stack. Such tools need only loose coupling and interoperation because layering Monte Carlo (MC) simulators is all one needs to do in order to study intricate correlations between parameters at various layers of the stack. However, MC simulations fail when one needs to describe highly-entangled quantum probability distributions, as MC techniques can not account for quantum interference between the different MC samples[2]. One needs either to (1) use approximate surrogate models of the lower layers (hence trivializing the noise model and losing track of important correlations) or to (2) create tools to convert between the many different ways quantum correlation might be internally represented by different special-purpose simulators.

In this position paper, we suggest the pursuit of a modular scalable quantum computing approach, focusing on building digital twin simulators and electronic design automation (EDA) for quantum computing with an emphasis on interoperable and automated physical-layer models.

Opportunity: Codesign leads to significant performance improvements (e.g. in entanglement purification [1,2]), but the quantum hardware design tools are lacking. On the other hand, today there is an unprecedented richness in reverse design tools, theory, and open source software, partially thanks to the growing investments into supporting technologies for AI. One particularly extraordinary example is the availability of automatic-differentiation compilers capable of providing high-performance gradient calculations over arbitrary programs, e.g., capable of differentiating even through programs with randomness (e.g. simulations involving wavefunction collapse). These autodiff technologies enable codesign and optimization that previously would have been prohibitively expensive. Moreover, ML

---

[1] corresponding author
[2] While the possibility for interference is one of the roots of the computational advantage inherent in quantum hardware, it is also the reason it is challenging to construct practically useful high-fidelity digital twins of quantum hardware as, e.g., both the low level analog physics simulator and the high level network simulator might need to track how they affect each other’s quantum correlations (not doable with classical Monte Carlo).

techniques have percolated into many other fields, providing unsurpassed optimization tools for a variety of quantum tasks[3,4]. Lastly, large language models have the potential to greatly automate the more menial research tasks, especially if we establish standardized formats for a more accessible scientific record. These features together have the potential to drastically speed up the development and optimization of quantum technologies if made available in the form of EDAs.

However, a tool to combine disparate simulators and marshal them into a synchronized super-simulator is needed (due to the unavailability of stacked MC) for high-fidelity digital twins. Simulating a quantum network to high fidelity requires both the modeling of the network dynamics[5], but also the low-level analog noise dynamics of each node, if we are to reap the benefits of codesign and optimization [1]. Thankfully, that does not require rewriting available simulator code, but it highlights the importance of building tools that can build large scale simulations by patching together and timing smaller special-purpose simulations, especially if we want to have good optimizers, not only simulators.

Assessment: The success of the proposed framework will be measured by its ability to efficiently design, analyze, and implement quantum algorithms, architectures, and networks, as well as its impact on the accessibility and extendability of scientific research. This includes assessing the performance of digital twin simulators and EDA tools in terms of model accuracy, interoperability, and ease of use, as well as the effectiveness of automated physical-layer models in their ability to streamline quantum computing research, facilitate machine verification, and promote interoperability among various quantum computing systems. A full-stack model that took a whole team and six months should become something that a single graduate student can build, study, and optimize in a few weeks.

A brief list of specific capabilities that are needed include (1) Support for various types of automatic differentiation and reverse design, probabilistic/Bayesian programming and parameter estimation, (2) Ease of creating numerical surrogates of real or simulated hardware, (3) Automating control of and model fitting to experiments, (4) Database of typical hardware properties and noise parameters, error correcting codes, common networking protocols and primitives, (5) Symbolic algebra system permitting formalism-and-simulator-agnostic representation of states and processes, (6) Superb debugging and visualization tooling, (7) Ease of publishing machine readable, interactive, reproducible research results.

Timeliness or maturity: The performance, scale, and sophistication of quantum hardware is at the precipice at which optimizing around the aforementioned bottlenecks would is for continuing progress. At the same time, the software building blocks to enable the creation of these tools only now became commonplace. We need these new tools in order to perform this design and optimization, and now we know how to build these tools.

References:

[1]: Krastanov et al. "Optically heralded entanglement of superconducting systems in quantum networks." Physical Review Letters 127.4 (2021): 040503.
[2]: Co-Design Center for Quantum Advantage. "Architectures for Multinode Superconducting Quantum Computers", arXiv:2212.06167.
[3]: Fösel et al. "Reinforcement learning with neural networks for quantum feedback." Physical Review X 8.3 (2018): 031084
[4]: Sivak et al. "Model-free quantum control with reinforcement learning." Physical Review X 12.1 (2022): 011059.
[5]: Coopmans et al. "Netsquid, a network simulator for quantum information using discrete events." Communications Physics 4.1 (2021): 164.

A Quantum Internet Beyond Dark Fiber: Coexistence of Quantum and Classical Light in the Optical Fiber Infrastructure for Scalable Quantum Networking

Jordan M. Thomas[1]*, Keshav Kapoor[2], Neil Sinclair[3], Joaquin Chung[4], Gregory S. Kanter[1], Prem Kumar[1]
[1]Northwestern University, [2]UIUC, [3]Caltech, [4]Argonne National Laboratory
email: jordanthomas2025@u.northwestern.edu, gregory.kanter@northwestern.edu, kumarp@northwestern.edu

Topics: *quantum network models and architectures; codesign and integration*

1 Challenges

As new quantum network technological capabilities develop, the extent to which such quantum systems can exist beyond laboratory settings and the economic cost of deploying them should be considered to evaluate the scalability of quantum networking in real-world settings. The most accessible medium for deploying quantum networks is within the optical fiber infrastructure, however most fibers are already populated with classical telecommunication systems and dark fibers are in demand with the expansion of classical communications. Therefore, one key challenge is to engineer quantum networks to "coexist" in the same fibers with classical signals as this will dictate the availability of fiber links and ease costs of leasing and maintenance of fibers [1]. Further, coexistence of classical and quantum light can be utilized for quantum network functions (e.g., time synchronization [2] and teleportation), where classical light encoding information about the quantum system coexisting in the fibers as quantum signals would simplify architecture designs and limited fiber link resources. Furthermore, the distribution of strong classical light is required to implement protocols based on continuous variables [3] or for single-photon protocols [4].

Such quantum-classical integration has been widely studied in the context of single-photon and single-channel quantum systems; the main challenge being spontaneous Raman scattering of high-power classical communications generating noise photons that can corrupt quantum channels. However, quantum network technology has been rapidly developing in recent years towards deploying significantly more complex quantum network systems (e.g., multi-node entanglement-based networks, quantum teleportation, entanglement swapping, and quantum memories) in dark fibers, however operating these systems beyond dark fibers remains largely understudied. The unique physics of each application leads to interesting new research directions for understanding how noise from coexisting classical signals will impact various components within future quantum network physical layers, which will then guide the codesign of scalable quantum networks not limited to dark fibers.

2 Opportunity

Completely integrating both quantum and classical communications in the same fibers will require multiple fields within physics, engineering, and networking in order to develop physical models for how noise impacts various quantum network physical components. Furthermore, we also need to engage theorists to understand how errors produced by noise permeate up the stack, all the way up to applications. We here argue that a research program that uses such models to engineer systems and control routing for noise robust quantum-classical network operations, control planes managing the routing of both quantum and classical signals, and error correction/mitigation strategies for applications is needed.

The main challenge for integration is noise from high power classical light, where a common method to reduce noise is to allocate quantum and classical light within different telecom bands. Thus, there may be situations wherein all telecom bands are populated with quantum and classical systems, and thus maturing technology across all telecom wavelengths should be pursued in both the quantum and classical domains.

- Advancing O-band classical transceivers across $1260 - 1300$ nm would reduce the noise in C- and L-band quantum networks and allow high data rate coexisting classical channels with minimal impact on

*Corresponding author: jordanthomas2025@u.northwestern.edu

quantum functions [5], and similar developments can be applied to the quantum domain using O-band quantum systems when C- or L-band classical communications are operating within a fiber [6].

- Classical and quantum network in the control plane perspective should cooperate to efficiently schedule access to the shared fiber infrastructure.
- Technology to achieve tight temporal filtering with high efficiency, i.e. an ideal detector could be superconducting nanowires with tight electronic or low-loss optical gating.

3 Assessment

Success can be evaluated by the ability to operate each quantum network function albeit the presence of high power classical signals operating in the same fibers. Mostly, this can be characterized by the fidelity of the quantum network operation as a function of the possible coexisting classical power levels that can be achieved before quantum protocols fail due to intolerable noise levels. As an example of the expected improvement, we estimated [7] the reduction in forward scattered noise for a fiber link using superconducting nanowires compared to conventional photodetectors used in previous work [2], finding orders of magnitude improvement, see Fig. 1.

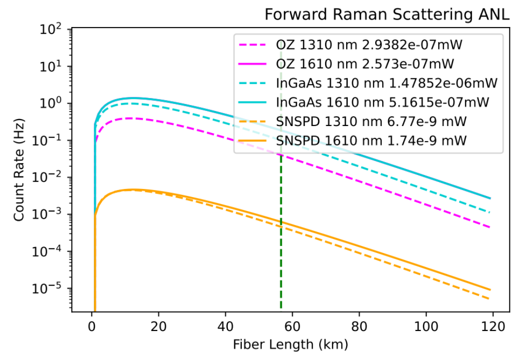

Figure 1: The measured count rate of forward Raman scattered light into the quantum channel C-band window are shown for 1610 nm and 1310 nm clock signals. The legend details the type of detector and launch power needed to reach the distance labelled on the horizontal axis. Calculations based on detectors used in ANL-FNAL fiber experiments [2].

4 Timeliness or maturity

Quantum network technology has begun to allow more complex quantum systems to be deployed within the fiber infrastructure. Recent work has demonstrated the feasibility of multi-channel entanglement-based networking alongside terabits per second classical communications in 48 km installed fiber with 18 dBm coexisting classical power [6]. As such quantum systems are key components in quantum network applications such as multi-node quantum networks, teleportation systems, and entanglement swapping, these methods can be readily built upon when moving towards these more complex quantum network applications coexisting alongside high data rate classical communications.

Successful integration of quantum and classical communications coexisting alongside each other in the same fiber infrastructure would render widespread deployment of quantum networks, and eventually a quantum internet, a more feasible reality.

References

[1] T.E. Chapuran, et. al., "Optical networking for quantum key distribution and quantum communications," New J. Phys. 11, 105001 (2009).

[2] K. Kapoor et al., "Picosecond Synchronization System for the Distribution of Photon Pairs through a Fiber Link between Fermilab and Argonne National Laboratories," IEEE J. Quantum Electron (2023).

[3] Jouguet, Paul, et al. "Experimental demonstration of long-distance continuous-variable quantum key distribution." Nature photonics 7.5 (2013): 378-381.

[4] Lucamarini, Marco, et al. "Overcoming the rate–distance limit of quantum key distribution without quantum repeaters." Nature 557.7705 (2018): 400-403.

[5] J. M. Thomas, et. al., "Optimization of Classical Light Wavelengths Coexisting with C-band Quantum Networks for Minimal Noise Impact," in Optical Fiber Communication Conference (OFC) 2023, p. Tu3I.3.

[6] J. M. Thomas, Gregory S. Kanter, Prem Kumar, "Designing Noise-Robust Quantum Networks Coexisting in the Classical Fiber Infrastcture," arXiv:2304.09076 [quant-ph] (2023).

[7] K. Kapoor, N. Sinclair, et al. in preparation (2023).

A Hybrid Free-space/Fiber Architecture for Long-distance Quantum Networks

Paul G. Kwiat[1,*], Makan Mohageg[2]
[1]University of Illinois at Urbana-Champaign, [2]NASA Jet Propulsion Laboratory
email: kwiat@illinois.edu

Topics: *Quantum network models and architectures; DOE Science applications*

1 Challenges

There is a large, and growing, effort to develop multi-node and long-distance quantum networks. These will enable new capabilities in communication, computing, and sensing. For example, with suitable communication rates and high fidelity, such links can be used to realize both secure remote quantum programming ("blind'' quantum computing), as well as distributed quantum processing, in which the already tremendous power of a quantum computer can be amplified exponentially by allowing multiple smaller processors to act coherently as one large machine [1]. As a rather different application, quantum networks can enable distributed quantum sensing, e.g., as a resource to as a resource of distributed quantum photonic states, which can be used to enable much longer very long baseline interferometry (VLBI), which in turn enables much higher resolution telescopy [2] (again, assuming sufficient rates and fidelity of the available quantum states).

Although there is little doubt that many of the quantum networking applications will rely on fiber-based communication channels, these do have the non-negligible disadvantage of a transmission that drops off exponentially over distance. Even working at the optimal telecommunication wavelength for fiber transmission (1550 nanometer), where the loss is only 0.2 dB/km), the net transmission over 100 km is only 1%, and the transmission over 800 kilometer (e.g., necessary to connect Fermi National Lab and Oak Ridge National Lab) is only 10^{-16}, while the transmission for a fiber link connecting Los Alamos National Laboratory and ORNL (2000 km) is 10^{-40}! Obviously, the problem gets exponentially worse when we consider transcontinental, or even transoceanic links.

2 Opportunity

One strategy to combat this problem is to employ quantum repeaters, nodes that employ small quantum processors, and multi-mode synchronizable heralded quantum memories. Unfortunately, such quantum technology is still under development, and today no single viable quantum repeater has been demonstrated which could run at rates relevant for distributed quantum processing. Given that the eventual spacing between such repeater nodes could be as low as 20 km [3], it will likely be quite some time before technology is at the level to enable transcontinental quantum entanglement distribution, or even between LANL and ORNL (which might then require $O[100]$ repeaters).

The other approach is to employ free-space optical channels to distribute quantum entanglement. For example, one could consider one or more satellites in a low earth orbit (similar to the International Space Station), used to distribute entanglement to widely separated ground stations [4]. In fact, just such a capability was demonstrated by the famous Chinese Micius satellite, which distributed entangled pairs to two ground stations separated by 1200 kilometers, at rates up to 10 orders of magnitude beyond what would have been achieved using direct transmission through fiber [4].

In any event, to be useful for most applications, the free-space links will need to connect to the fiber-based quantum channels, since those are how the intended users, e.g., national labs, will be connected. For this reason we believe it is important to explore hybrid quantum network architectures, combining fiber-based and free-space links.

Some of the same advantages can be realized using terrestrial free-space links, which have the added advantage (over fiber) that they can be readily established and reconfigured without laying new optical

*Corresponding author: kwiat@illinois.edu

fibers. Furthermore, free-space links can enable connections over networks for which fiber connections may be essentially impossible, i.e., to moving vehicles.

Another advantage to space-based quantum networking is the possibility to explore new science. For example, the much higher velocities, enable one to study quantum nonlocality in moving inertial frames. Perhaps even more exciting is the opportunity to look at the effects of quantum measurement when gravitational time dilation is resolvable [5].

3 Assessment

Many challenges still remain, however. For example, because the loss increases quadratically over distance due to diffraction, it is necessary to have fairly large transmitting and receiving telescopes, which then necessitates the need for adaptive optics to be able to couple the light into single-mode optical fibers (where it can be used to realize entanglement swapping with other terrestrially-based entanglement sources). In fact, to date there have been no experimental demonstrations at all that have achieved entanglement swapping, or even the simpler cousin, quantum teleportation, between two sources in relative motion, which of course would be necessary for any space-based implementation. The principal challenge here is synchronization—the requisite Bell-state analysis requires the two photons to meet on a beamsplitter simultaneously (to within their coherence time, which varies inversely as the photon bandwidth, and in nearly every teleportation experiment to date is 0.1 -10 ps).

Specific areas where development are necessary to enable high-rate high-fidelity free-space and space-based entanglement networking include:

- Development of optimized low-SWaP source of 'swappable' entangled photon pairs; the requirement for swappability means the photons need to be indistinguishable, and not have unwanted spectral correlations. Such sources should also be multiplexable, to allow the necessary transmission rates;
- Development of high-efficiency and tunable frequency conversion; this is critical, both to allow the photons to be able to match available quantum memory architectures, but also to function in the presence of sizable Doppler shifts (e.g., a telecommunication-wavelength photon transmitted from/received by a LEO satellite can see Doppler shifts as high as 20 GHz);
- Development of suitable synchronization methods to enable entanglement swapping between sources in relative motion. Note that such solutions are also necessary for purely terrestrial network connections, due to drifts in fiber channel lengths, variations during routing via various links, etc.
- Development of flyable high-efficiency low-jitter detectors, e.g., superconducting nanowire single-photon counters, or Si avalanche photodiodes plus frequency upconversion if necessary.

4 Timeliness or maturity

Given the recent satellite entanglement distribution, and free-space quantum communication demonstrations, now is the appropriate time to begin serious research into free-space and space-based architectures, as these will undoubtedly take some number of years to bring to maturity. Nevertheless, the methods and technologies developed will likely benefit traditional fiber-based quantum networking as well.

References

[1] Rodney Van Meter and Simon J. Devitt. The Path to Scalable Distributed Quantum Computing. Computer, 49(9): 31–42, September 2016. ISSN 0018-9162. doi: 10.1109/MC.2016.291.

[2] Daniel Gottesman, Thomas Jennewein, and Sarah Croke. Longer-Baseline Telescopes Using Quantum Repeaters, Phys. Rev. Lett., 109(7), 070503, 2012, doi: 10.1103/PhysRevLett.109.070503.

[3] Mihir Pant, Hari Krovi, Dirk Englund, and Saikat Guga. Rate-distance tradeoff and resource costs for all-optical quantum repeaters, Phys. Rev. A, 95(1), 2017, doi: 10.1103/PhysRevA.95.012304.

[4] Ji-Gang Ren et. al. Ground-to-satellite quantum teleportation, Nature 70(73), 2017, https://doi.org/10.1038/nature23675

[5] Makan Mohageg et. al. The deep space quantum link: prospective fundamental physics experiments using long-baseline quantum optics, EPJ Quantum Technology, 9(1), 2022, https://doi.org/10.1140/epjqt/s40507-022-00143-0

Prospects for interconnecting quantum networks

Ying-Cheng Lai,[1,†] Houlong Zhuang,[1] and Joseph M. Lukens[1,2]

[1] *Arizona State University, Tempe, Arizona 85287, USA*
[2] *Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA*
[†] *ying-cheng.lai@asu.edu*

Abstract: Like its classical forerunner, the defining characteristic of the anticipated quantum internet will be the interconnection of many smaller quantum networks together. The development of a successful quantum internet will therefore require solutions to difficulties that, although related, are distinct from those faced by smaller testbeds. In this position paper, we describe important challenges and opportunities toward the realization of quantum internetworking that together form necessary—if not sufficient in themselves—conditions toward a truly quantum internet.

Topics: applications, models, codesign and integration.

*Challenge.—*The fundamental challenges of quantum internetworking arguably all stem from the no-cloning theorem: arbitrary quantum states cannot be copied without degrading their fidelity [1]. Consequently, many of the technologies that have proven so successful in the classical internet—e.g., optical amplifiers and electronic routers—are unavailable to quantum signals. At a basic level, some of the most conspicuous unmet research needs pertain to the hardware required for interconnected quantum devices, which can feed into a network hierarchy modeled after the classical internet: namely, quantum local area networks (QLANs) connected to form quantum metro area networks (QMANs), and QMANs connected to form quantum wide-area networks (QWANS), of which the quantum internet is the ultimate example [2,3]. Figure 1(a) depicts schematically how a single QLAN could connect to additional quantum networks through this structure, along with the fundamental quantum building blocks an internetwork would comprise: nodes, switches, routers, and repeaters. Underneath the visual simplicity of the icons in Fig. 1(a) lies extensive technical challenges. For with the exception of point-to-point quantum key distribution (QKD) systems, there currently exist no versions of these elements with the type of plug-and-play maturity expected from their classical precursors.

*Opportunity.—*Promising opportunities to address these challenges are emerging, however. As a specific example, wavelength-selective switches (WSSs) have been introduced to quantum networking for reconfigurable control of quantum spectral channels on demand [4]. Leveraged for entanglement distribution in a QLAN [5], WSSs can be nested as highlighted in Fig. 1(b); in this picture focused on entanglement distribution, a single WSS supports intra-QLAN entanglement, while nested WSSs facilitate inter-QLAN connections, analogous to switches and routers, respectively, in Fig. 1(a). Such an analogy should not be taken too far, however, as WSS-based internetworking in this model does not offer the packet switching functionalities expected from conventional electronic switches and routers. Whether utilizing WSSs or not, an important unfulfilled internetworking research opportunity will be to demonstrate on-demand entanglement between nodes from QLANs that are truly distinct, as defined by self-sufficiency and independence (as is the case with classical LANs).

For a future QWAN to earn the moniker "quantum internet," it will certainly need to be scalable, in the sense that smaller quantum networks can continuously be added and removed without any disruptions to its basic service. One of the key technological enablers of the current internet's scalability is packet switching, in which messages are routed through the network core on the fly through storage and forwarding, rather than via a preallocated communication pathway in circuit switching. Although all-optical packet switching has been studied in conventional networks, its cost and complexity have precluded serious competition with digital electronics. It is possible that quantum internetworking could finally provide the economic and scientific weight needed to revisit all-optical packet switching on a large scale; recent analyses of QKD packet switching [6] might be a harbinger for research to come.

An additional but perhaps underappreciated requirement of the quantum internet will be global timing synchronization; coherent operations and measurements on entangled quantum systems demand tight synchronization that can easily fall in the deep sub-ns range. Such low jitters are orders of magnitude below the ms- or µs-levels attainable with the Network or Precision Time Protocols (NTP or PTP), respectively, common in Transmission Control Protocol/Internet Protocol (TCP/IP) networks. Auspiciously, White Rabbit—a suite of protocols and hardware developed at CERN—satisfies many of these demands remarkably well, operating over optical fiber and attaining ps-scale jitters [7], qualities which have been exploited for significant reductions in noise in a deployed QLAN [8]. Yet although White Rabbit is inherently scalable in terms of nodes, its performance is distance- and hardware-dependent, pointing to the need for substantial engineering if this approach is to be scaled to internet levels. Additionally, it is likely that coexis-

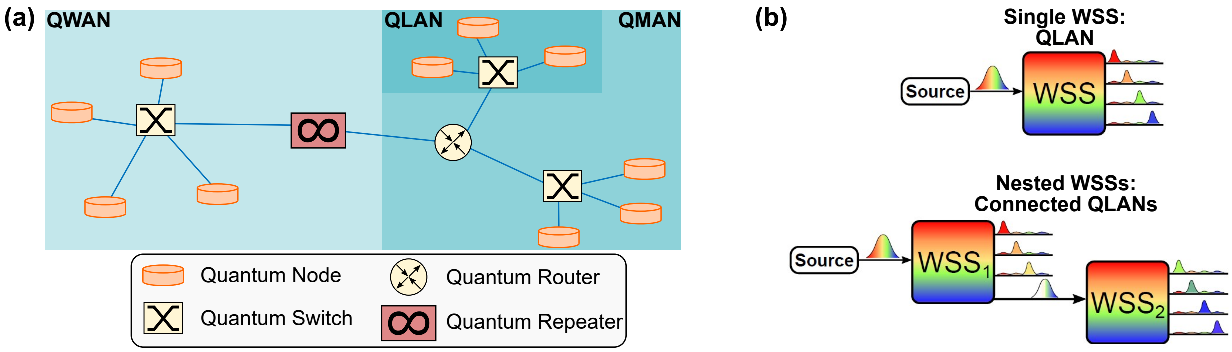

Fig. 1. (a) Representative components of a quantum internet, comprising QLANs connected to form QMANs, and QMANs connected to form QWANs. (b) Example physical topology where WSSs facilitate entanglement both within and between distinct QLANs.

tence of classical and quantum signals in the same optical fiber will prove a necessity for scalable QWANs. Whereas access to dedicated dark fiber might be attainable for QLANs self-contained within individual organizations, the cost associated with dark fiber links spanning transcontinental or transoceanic distance suggests much stronger pressure for coexistence in quantum *internetworking* than quantum *networking* as such. Many questions about performance for quantum applications beyond QKD remain; incidentally, preliminary studies of the noise generated by White Rabbit signals [9] show promise for attaining both synchronization and coexistence simultaneously.

Assessment.—Even after successful deployment, interconnected quantum networks will need to be evaluated and subsequently managed. In the flex-grid example of Fig. 1, such management includes provisioning bandwidth to best meet user demand according to some metric. Models and optimization approaches are progressing toward this goal, but so far have only been analyzed in the single-WSS (QLAN) case, where light paths follow a simple star (hub and spoke) topology [10]. Much more general global models and metrics that can efficiently manage quantum internetworks require ongoing research. Along these lines, codevelopment of management procedures with the physical quantum hardware is critical to ensure practicality. While tempting to seize upon the success of the TCP/IP stack as motivation for abstraction in quantum network research now, the current immaturity of quantum sources and processors makes it unclear whether a given abstraction will remain consistent with experimental validation. For these reasons, recent tests of quantum stacks in specific experiments provide a valuable template [11]. To enable the level of abstraction ultimately required for the quantum internet, an organic development of the quantum stack by physicists and network engineers will be required. Is the TCP/IP stack a useful guide for the quantum internet, or an artificial burden? Answering questions like this will be vital for assessing the success of quantum internetworking theory and experiment.

Timeliness.—The success of the existing classical internet, coupled with the promise of distributed quantum information processing, has led to an explosion of recent interest in the development of the quantum internet [12], including from DOE specifically [2, 3]. Although quantum networking as a field remains in a relatively nascent phase, rapid progress points to the critical importance of quantum internetworking research now, in order to support interoperability and prepare for the connection of multiple smaller testbeds in the near future.

References

1. W. K. Wootters and W. H. Zurek, Nature **299**, 802 (1982).
2. U.S. Department of Energy, “Quantum networks for open science” (2019).
3. U.S. Deparment of Energy, “From long-distance entanglement to building a nationwide quantum internet: report of the DOE Quantum Internet Blueprint workshop” (2020).
4. N. B. Lingaraju, H.-H. Lu, S. Seshadri, D. E. Leaird, A. M. Weiner, and J. M. Lukens, Optica **8**, 329 (2021).
5. M. Alshowkan, B. P. Williams, P. G. Evans, N. S. V. Rao, E. M. Simmerman, H.-H. Lu, N. B. Lingaraju, A. M. Weiner, C. E. Marvinney, Y.-Y. Pai, B. J. Lawrie, N. A. Peters, and J. M. Lukens, PRX Quantum **2**, 040304 (2021).
6. R. Mandil, S. DiAdamo, B. Qi, and A. Shabani, arXiv:2302.14005 (2023).
7. M. Rizzi, M. Lipinski, P. Ferrari, S. Rinaldi, and A. Flammini, IEEE Trans. Ultrason. Ferroelectr. Freq. Control. **65**, 1726 (2018).
8. M. Alshowkan, P. G. Evans, B. P. Williams, N. S. V. Rao, C. E. Marvinney, Y.-Y. Pai, B. J. Lawrie, N. A. Peters, and J. M. Lukens, J. Opt. Commun. Netw. **14**, 493 (2022).
9. I. A. Burenkov, A. Semionov, Hala, T. Gerrits, A. Rahmouni, DJ Anand, Y.-S. Li-Baboud, O. Slattery, A. Battou, and S. V. Polyakov, Opt. Express **31**, 11431 (2023).
10. J. Alnas, M. Alshowkan, N. S. V. Rao, N. A. Peters, and J. M. Lukens, Opt. Express **30**, 24375 (2022).
11. M. Pompili, C. Delle Donne, I. te Raa, B. van der Vecht, M. Skrzypczyk, G. Ferreira, L. de Kluijver, A. J. Stolk, S. L. N. Hermans, P. Pawełczak, W. Kozlowski, R. Hanson, and S. Wehner, npj Quantum Inf. **8**, 121 (2022).
12. S. Wehner, D. Elkouss, and R. Hanson, Science **362**, eaam9288 (2018).

Assessing the potential of quantum advantage in quantum chemistry with state-of-the-art quantum algorithms

Joonho Lee
Assistant Professor, Department of Chemistry and Chemical Biology, Harvard University,
joonholee@g.harvard.edu

Topic: Algorithms – (a) quantum algorithms admitting theoretical or empirical evidence of advantage for fundamental domains such as simulation, optimization, or machine learning, (b) hybrid quantum and classical algorithms

Challenges: Quantum phase estimation (QPE) and many other quantum algorithms have been put forward with the expectation that these will demonstrate a quantum advantage in quantum chemistry for some ground-state applications. However, no concrete evidence exists for a practical quantum chemistry application with a speedup with any existing quantum algorithms.

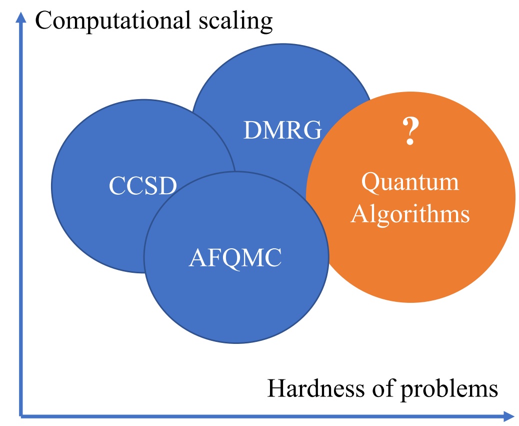

Figure 1 Potential phase diagram with quantum advantage compared to popular classical (CCSD, DMRG, AFQMC) methods.

An example algorithm is the quantum-classical hybrid quantum Monte Carlo (QC-QMC) algorithm developed by our group [1]. In Ref. [1], we implemented QC-QMC on Google's Sycamore processor and scaled the algorithm up to 16 qubits despite the presence of circuit noise. Our QC-QMC demonstration has been the largest quantum simulation of chemistry on an actual quantum computer to date. Despite its success, a potential quantum advantage through QC-QMC is also not obvious, as this depends on systems, quantum circuits, and QMC sampling [2].

Another example is our recent work in collaboration with the Chan group at Caltech, where we examined the evidence of quantum advantages in using QPE for challenging quantum chemistry problems such as FeMoCo [3]. Due to the cost of the initial state preparation, we concluded that it is unclear whether QPE will achieve a speedup over other known classical methods for classically challenging problems.

Assessing quantum speedups in quantum chemistry requires careful analysis of *both* classical and quantum algorithms (See Figure 1). Moreover, it depends on the precise details of the problem of interest, so one has to evaluate individual instances for a quantum advantage. These pose significant challenges in establishing a clear classical-quantum boundary in quantum chemistry.

Opportunity: Going beyond our recent assessment of QPE, there remain exciting opportunities for examining the prospects of quantum advantage in ground-state applications.

For QPE, one can efficiently prepare a low bond-dimension matrix product state (MPS) initial state [4]. Although there is a heuristic aspect in ensuring the qualitative correctness of a low bond-dimension MPS state, this offers a new, unexplored way to avoid the state preparation overhead that a simple mean-field initial state would incur. Our group is examining this exciting possibility, specifically in the context of simulating models of metalloenzymes.

We are also evaluating the practicality of adaptive variational quantum algorithms (i.e., ADAPT-VQE) [5]. ADAPT-VQE is one of the more widely used and discussed near-term algorithms in chemistry. Despite promising results on small systems, their full potential has never been examined because the corresponding quantum circuits are too expensive to emulate classically. To overcome this challenge, our group has implemented an adaptive version of the Hartree-Fock theory (ADAPT-HF) that runs efficiently on a classical computer. We are examining the practicality of the ADAPT algorithm at scale with this simplest possible ansatz.

Next, because QC-QMC avoids optimization, one can design a set of numerical experiments to probe the precise quantum-classical boundary for QC-QMC without the danger of local minima. We recently established a clear limit for classical auxiliary-field QMC (AFQMC) [6], which can handle any binuclear transition metal complexes with near-exact accuracy. We are now performing a numerical investigation of binuclear complexes to understand how small each statistical overlap sample becomes during AFQMC. These overlap values determine the measurement overhead in QC-QMC. We hope to establish the boundary of quantum and classical computing within QC-QMC for this set of challenging chemical systems. Upon establishing such a boundary, we will run larger QC-QMC calculations on a quantum computer (e.g., Sycamore).

Finally, we recently proposed real-time electronic dynamics [7] as a potential venue for observing quantum advantage. Unlike the previously mentioned examples, these applications circumvent the need for accurate state preparation and the exponential-scaling overhead due to that. We currently plan to study the quench dynamics in prototypical two-dimensional lattice systems, such as the Heisenberg model, with the aim of examining the limits of classical and quantum algorithms. Our investigation will lead to co-designing an experiment on the Lukin (Harvard) group's Rydberg atom array.

Assessment: In all four algorithm use cases, we strive to define the elusive boundary between quantum and classical algorithms. While such a boundary is application-specific and algorithm-specific, the success of these research directions will lead us to specific use cases of quantum computers with an expected quantum advantage. In the case of QC-QMC and real-time dynamics, we will co-design a specific set of experiments near the boundary to demonstrate one of the first practical quantum advantage in chemistry and materials science. For a quantum advantage prior to full error correction, we expect 50 physical qubits or more to be needed for both QC-QMC and real-time dynamics. Quantum computers of this size are readily available today with superconducting qubits or Rydberg atom arrays.

Timeliness or maturity: We have many quantum algorithms developed for quantum chemistry since the first QPE proposal for quantum chemistry. Setting aside the circuit noise, we started to assess the prospects of quantum algorithms such as QPE. To move forward as a field, it is critically important to evaluate the evidence for a quantum advantage in each of the available algorithms. We have new promising algorithms, such as QC-QMC, and a new research direction beyond ground state, such as real-time dynamics, where one may find quantum advantage prior to error correction. The need for assessing evidence for these new algorithms is at its peak, based on our recent findings. Even if we fail to find a quantum advantage before error correction, having a clear quantum-classical boundary will be of extensive interest and value in this field.

References:

[1] Huggins, …, **Lee**, *Nature* **603**, 416 (2022).
[2] **Lee**, …, arXiv:2207.13776.
[3] Lee, **Lee**, et al., *Nat. Comm.* **14**, 1952 (2023).
[4] Rudolph et al., arXiv:2209.00595.
[5] Grimsley et al., *Nat. Comm.* **10**, 3007 (2019).
[6] **Lee** et al., *JCTC* 18, 7024 (2022).
[7] Babbush, …, **Lee**, *Nat. Comm.* (Just accepted; arXiv:2301.01203).

On the Need for Extensible Quantum Compilers with Verification

Tyler LeBlond, Xiao Xiao, Eugene Dumitrescu, and Ryan Bennink
Oak Ridge National Laboratory

Alexandru Paler
Aalto University

I. TOPIC

Key-words: Compilation, error correction and mitigation, and codesign and integration.

In this position paper, we posit that a major DOE-funded open-source quantum compilation platform is needed to facilitate (a) resource optimization at the fault-tolerant layer of the quantum computing software stack and (b) co-design of that layer of the stack with other layers, and that this platform needs to be extensible and include verification.

II. CHALLENGE

The surface code is the most popular quantum error-correcting code (QECC) due to its high error threshold and nearest-neighbor connectivity, which improves its applicability to real hardware [1]. It is therefore important for the broader community to focus on developing surface code compilation frameworks, and several authors have proposed such frameworks (either abstractly or in software) using so-called lattice surgery, which is the leading technique for implementing fault-tolerant operations with surface codes [2–8]. A review of these works reveals that there are multiple layers of complexity to the development of a surface code compilation framework. One must (at least) choose (a) an input logical circuit instruction set, (b) an output surface code instruction set, (c) a physical layout of surface code patches that includes an appropriate amount of ancillary space for transport and resource state generation, (d) an abstract way to re-express the input logical circuit into surface code operations, and (e) a scheme for scheduling, routing, and further breaking down the list of abstract surface code instructions into something physically realizable. Because fault-tolerance is expected to increase the overall space-time cost of quantum computation by several orders of magnitude, a fuller optimization of this layer of the stack (b, c, and d) together with its co-design and integration with other layers (a and e) is an important present challenge for the realization of quantum advantage in the fault-tolerant era.

While open-source software packages for surface code compilation have been published, e.g. [7, 8], there is no consensus on which compilation strategy will perform the best on practical quantum circuits. Moreover, we must ask the question: even if a robust lattice surgery compiler is written, what happens when a more resource-efficient fault-tolerant implementation of logical gates on the surface code is proposed? Despite recent success in implementing the surface code, there is reason to believe that it might not remain the preferred QECC long-term. Its encoding rate does not scale favorably compared with other members of the low-density-parity-check (LDPC) code family (these can encode increasingly many logical qubits with an asymptotically linear number of physical qubits [9], which is a very desirable property for resource efficiency reasons). An understanding of how to implement LDPC logical gates fault-tolerantly and resource-efficiently is a subject of ongoing research, though it is suspected that lattice surgery protocols could apply to them. With these realities on the horizon, future quantum compilers should have extensibility to broader classes of QECC in mind.

Additionally, because executed quantum computations have to be both fault-tolerant and correct at the same time, it is critical that the QECC-compiled circuits are verified by the compiler. Several things could go wrong in the process: there might be software bugs in the compiler, the compilation/optimisation methods might be formally incorrect, or the input circuit might even be wrong. However, considering that more general LDPC codes might form the foundation of future large scale computers, it is necessary to devise compilers (or layers of the compilation stack) that can offer guarantees about the compiled circuit. For instance, are the logical gates implemented correctly by the surgery protocols? Does the resulting gate sequence reflect the input circuit? Verification can be performed in approximate [10] or exact manners [11]. The latter is exponentially difficult to perform at scale (state vector simulation is very expensive, tensor network approaches are a bit less expensive) while the first offers only bounded guarantees. Nevertheless, approximate verification (i.e. testing) is the only practically viable option.

Because there have been few attempts to create fault-tolerant compilers with verification, doing so represents a research frontier. While verification can be performed at different layers of the quantum computing stack (in bottom-up manner: physical qubit operations, fault-tolerant protocols, logical gates, higher-level logical constructs, functional verification of arithmetic units, . . . , oracles), efficient and scalable fault-tolerant compilers should at least support verification at the layers of the stack that they touch i.e. the logical compilation to fault-tolerant operations and the lower-level implementation of those into stabilizer measurement circuits.

III. OPPORTUNITY

With these considerations in mind, it appears to us that a platform is needed to compare the performance of different quantum compilation strategies on practical circuits and that this platform should be both *extensible* and *verified*. We expect that the level of effort required to develop this platform is significant enough that it is currently only amenable to companies with large, well-funded quantum software divisions. In this setting there is more incentive to produce quantum compilers that are highly tailored toward companies' specific aims and less incentive to produce an open-source software platform for the community as a whole to explore the vast space of possible choices on the different compilation layers. Therefore, we believe that there is an opportunity for a public entity such as the DOE to fund this platform and drive progress in quantum compilation.

IV. ASSESSMENT

The best overall outcome would be an open-source quantum compilation platform that forms the basis for future compilers written by companies in the private sector. Near-term, it should produce a benchmarking effort that establishes the performance of different surface code compilation frameworks on practical quantum circuits. This can help guide the development and scaling of fault-tolerant quantum computers. Long-term, it should be used as a sandbox for the development of fault-tolerant implementations using different QECC.

V. TIMELINES AND MATURITY

Since the advent of noisy intermediate-scale quantum (NISQ) devices, there has been an immense amount of literature published around the topics of QECC and fault-tolerant logical operations using them. This surge has resulted from the need to prepare for a coming fault-tolerant era, which has now been ushered in through initial demonstrations of fault-tolerance by major hardware providers [12].

---

[1] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Physical Review A **86**, 032324 (2012).
[2] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, New Journal of Physics **14**, 123011 (2012).
[3] A. G. Fowler and C. Gidney, Low overhead quantum computation using lattice surgery, arXiv preprint arXiv:1808.06709 (2018).
[4] D. Litinski, A game of surface codes: Large-scale quantum computing with lattice surgery, Quantum **3**, 128 (2019).
[5] M. Beverland, V. Kliuchnikov, and E. Schoute, Surface code compilation via edge-disjoint paths, PRX Quantum **3**, 020342 (2022).
[6] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoeffler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vaschillo, Assessing requirements to practical quantum advantage, arXiv preprint arXiv:2211.07629 (2022).
[7] A. Paler and A. G. Fowler, Opensurgery for topological assemblies, in *2020 IEEE Globecom Workshops (GC Wkshps* (IEEE, 2020) pp. 1–4.
[8] G. Watkins, H. M. Nguyen, V. Seshadri, K. Watkins, S. Pearce, H.-K. Lau, and A. Paler, A high performance compiler for very large scale surface code computations, arXiv preprint arXiv:2302.02459 (2023).
[9] N. P. Breuckmann and J. N. Eberhardt, Quantum low-density parity-check codes, PRX Quantum **2**, 040101 (2021).
[10] A. Paler and S. J. Devitt, Specification format and a verification method of fault-tolerant quantum circuits, Physical Review A **98**, 022302 (2018).
[11] Q. Wang, R. Li, and M. Ying, Equivalence checking of sequential quantum circuits, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **41**, 3143 (2021).
[12] Suppressing quantum errors by scaling a surface code logical qubit, Nature **614**, 676 (2023).

Computer-Aided Assessment of Early Fault-Tolerant Quantum Applications Towards Practical Quantum Advantage

Hongye Hu, Harvard Quantum Initiative Postdoc Fellow, Harvard University
Gushu Li, Assistant Professor, University of Pennsylvania
Contact: hongyehu@g.harvard.edu, gushuli@seas.upenn.edu (participating author)
Topics: algorithms, applications, and error correction and mitigation

Challenge

Quantum applications have been traditionally designed for either noisy intermediate-size quantum (NISQ) devices without quantum error correction (QEC) or hypothetical large-scale fault-tolerant quantum computers with full quantum error correction capabilities. While the demonstration of "quantum supremacy" on NISQ devices has been a major milestone, it remains uncertain whether such machines can provide practical advantages. On the other hand, the realization of a fully fault-tolerant quantum computer is still a distant goal due to its very high hardware resource requirement.

Recently, quantum processors have been significantly advanced, and QEC protocols have been experimentally demonstrated [1, 2, 3], which can significantly reduce the logical error rate and decoherence. The transition from NISQ to fault-tolerant quantum computing has been initiated, and we are entering the era of early fault-tolerant quantum computing. As more qubits are integrated into early fault-tolerant quantum processors and lower error rates are provided by QEC, there is more potential to demonstrate practical quantum advantages on these devices. However, assessing the quantum advantage of applications that can be experimentally demonstrated on such early fault-tolerant quantum computers presents significant challenges, which can be summarized as follows:

- **Application Instance Selection** There exist many quantum algorithms/applications and the programs of each quantum application may also vary with different input datasets, different targeted problems, etc. These quantum program instances will have different performances on the quantum computers with early fault-tolerant capability due to their built-in program structure, noise resilience, the QEC protocol configuration, etc. *It is not clear what quantum application can benefit more from such quantum computers and thus has the most potential to demonstrate practical quantum advantage.*
- **Ensuring Correctness/Results Quality** Assessing quantum applications executed on such early fault-tolerant quantum computers becomes increasingly difficult. The quantum processors are employing an ever-increasing number of physical qubits to accommodate the required redundancy in QEC and provide the potential to demonstrate practical advantages. However, it becomes intractable to simulate quantum applications with such many qubits on even classical supercomputers, making it *hard to verify the results of these quantum applications and assess their advantage over classical computing*.

These challenges require innovative approaches to assess quantum applications that can take advantage of early fault-tolerant quantum computers before establishing practical quantum advantages on these quantum devices.

Opportunity

We identify that computer-aided tools can be employed to tackle the challenges in assessing early fault-tolerant quantum applications. Some of the opportunities are listed as follows:

- **Formal Method** To assess the quality of the execution results, formal methods can be used to verify the program specification and quantify the execution time and accuracy for quantum programs executed on early fault-tolerant quantum computers. Without actually simulating the program, formal methods can provide mathematical guarantees for the correctness of the program. The result accuracy can also be obtained after incorporating the early faulty-tolerant quantum hardware characteristics. When the correctness of the application is specified by abstracted properties, more efficient formal verification is possible.

- **Synthetic Benchmark** For quantum applications with practical advantages, it should be hard to directly obtain the correct results using classical computers. But in the opposite direction, once we have an answer, it is possible to construct a quantum application instance with such a preset answer. Benchmarks for early fault-tolerant quantum computers can be automatically generated in this way by computer-aided tools. Then the correctness of the quantum applications can be tested by checking if the actual execution results from early fault-tolerant machines match the preset results.
- **Finding Physical Experimental Configuration** For some applications, especially quantum simulation, physical experimental validation is a possible solution. For example, the bond vibration frequency of a molecule can be experimentally measured and compared with the results obtained from the quantum simulation. Computer-aided tools can help find quantum program instances that can be accommodated by existing/near-term physical experiments. Such tools can be developed upon collecting existing experimental configuration/result data. Machine learning models can then be trained to predict the problem instances that can be experimentally verified.
- **Projected Classical Resource Estimation** Regarding the quantum advantage, the advantage of early fault-tolerant quantum applications can be evaluated by comparing their performance with their classical correspondence on supercomputers. This comparison can be made by developing analytical models and resource estimators to estimate and project the computational resources required by the best-known classical algorithm to solve the same problem. Finally, we can use the assessment results to refine the design of the applications and further improve their performance.

These methods can be used individually for their targeted task, or combined to systematically search for assess quantum applications with potential practical advantages. Collaborations among mathematicians, physicists, and computer scientists can naturally be triggered when investigating these opportunities.

Assessment

The topic of this white paper is to develop computer-aided approaches to assess the computational advantage of early fault-tolerant quantum applications. The central metric is whether the proposed tool can find such an application instance whose quantum advantage can be finally experimentally demonstrated. For the formal methods, the key metric is the size of the quantum application instance whose non-trivial properties can be checked in an acceptable time. The synthetic benchmark generation tool can be evaluated by generating small application instances that lie in the computation capability of existing quantum hardware. For the machine models that can predict if the application can be experimentally verified, the key metric is prediction accuracy. The actual number of problem instances that can be experimentally verified as predicted by the model is also important. After considering all possible optimizations.

Timeliness & Maturity

It is a timely topic to assess early fault-tolerant quantum applications with practical advantages. In the past, this has been challenging due to limited control of qubits and insufficient qubit counts. However, recent advancements in quantum hardware have provided processors with a sufficient number of qubits that can provide the necessary redundancy to realize early fault tolerance. Moreover, advancements in the peripheral control hardware allow fast feedback control to a massive array of physical qubits to provide QEC capability. Successfully executing this research will deliver quantum applications with practical advantages that can be demonstrated on near-future hardware platforms with early fault tolerance and provide insights for the future transition from early-fault tolerance to quantum computing with full fault tolerance.

References

[1] Google Quantum AI. Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 2023.
[2] S. Krinner et al. Realizing repeated quantum error correction in a distance-three surface code. *Nature*, 2022.
[3] V. V. Sivak et al. Real-time quantum error correction beyond break-even. *Nature*, 2023.

The Next Bottleneck: Quantum Control for Large-Scale Quantum Processors

Benjamin Lienhard (blienhard@princeton.edu),[1, 2] Robert L. Kosut (kosut@scsolutions.com),[3, 1]
Gaurav Bhole (gaurav.bhole@princeton.edu),[1] and Herschel A. Rabitz (hrabitz@princeton.edu)[1]
[1] *Department of Chemistry, Princeton University, Princeton, NJ 08544, USA*
[2] *Department of Electrical and Computer Engineering, Princeton University, Princeton, NJ 08544, USA*
[3] *SC Solutions, Sunnyvale, CA 94085, USA*
(Dated: May 1, 2023)

Keywords: Enabling viable **error correction:** Scalable QIP characterization, calibration, & stabilization

Quantum computers promise to solve specific computational problems significantly more efficient than contemporary computers [1]. However, before universal gate-based quantum computers [2] can become a reality, several more milestones must be reached and engineering challenges overcome. With the achievement of quantum supremacy [3, 4], error-corrected quantum computing is the next declared milestone towards practical quantum computers. While small-scale quantum information processors perform on a level sufficient for the most basic quantum error correction protocols [5], the control performance generally degrades rapidly as the number of qubits increases due to often uncharacterized system nonidealities such as crosstalk [6, 7] or unstabilized system parameter fluctuations [8, 9]. The increasing strength and number of nonidealities as quantum processors increase in size require immediate attention to enable the successful implementation of quantum error correction protocols at scale and the execution of useful quantum algorithms.

Realizing the dream of quantum computing is contingent on developing scalable, efficient, and robust quantum processor characterization, calibration, and operational stabilization routines calling for a dedicated and targeted funding program.

Today's quantum processor calibration methods do generally not address **crosstalk** beyond two qubits as control signals are independently characterized and calibrated for each control gate [3]. System crosstalk can occur due to control signal interference, residual coupling between control signals and qubits, or classical and quantum interactions between qubits, leading to undesired classical or quantum correlations. Uncontrolled correlations are detrimental to the realization of effective quantum error correction codes. Crosstalk is often neglected or compensated to first-order using on-demand coupling schemes [3, 10]. **To address the challenge of crosstalk, we need to develop *global characterization tools* that remain efficient as systems scale.**
The **calibration of control pulses** typically starts with a system-dependent ansatz that is subsequently fine-tuned by monitoring the response of each qubit [3]. Automated protocols using dedicated classical hardware [11, 12] have recently superseded manual pulse optimization. To further reduce the effort, elaborate calibration algorithms have been conceived [13]. First attempts have been made to employ cloud-based reinforcement learning—powerful but notoriously computationally intensive machine learning algorithms—to improve the calibration performance [14]. While these methods can support tens to maybe hundreds of qubits, the needed computational power will rapidly exceed what is available through existing supercomputers and prohibit passing the targeted milestone of error correction. **To address the challenge of calibrating thousands to millions of qubits, we need to investigate the *potential and limitations of classical resources* and *develop efficient closed-loop calibration algorithms* that leverage the optimization landscape [15] and exploit the typically incomplete knowledge of the system dynamics and error models.**
Lastly, system parameters, such as the transition frequency of a qubit or the coupling to its environment, are often subject to **temporal fluctuations** [16]. This temporal **system parameter drift**—discrete and continuous [9]—results in a degradation of the control performance and calls for periodic and laborious re-calibration. While decoupling sequences can suppress dephasing, they do not address the system parameter drift. **To reduce the need for repeated calibration and improve the quantum processor performance, we need to develop *predictive operational stabilization routines* [17].**

To enable the successful implementation of viable quantum error correction protocols at scale, we recommend exhausting the capabilities of quantum control and its tools first. To do so, we need to develop characterization tools to identify global system nonidealities, establish the full potential of classical resources, develop calibration algorithms lean in data demands, and stabilize operational system parameters in a predictive manner, elaborated in the table below.

Challenge	Opportunity	Impact	Time
Classical & Quantum Crosstalk	**Global characterization to identify crosstalk beyond nearest neighbors**	Identifying correlations beyond nearest neighbors will enable the development of targeted quantum control routines to suppress the correlations and thus reduce their detrimental effects on effective quantum error correction.	< 2028
Large-Scale Calibration	**Establish and explore the full capabilities of classical resources**	Establishing their full potential will define the playground for future calibration routines.	< 2027
	Efficient closed-loop calibration algorithms	Equiping closed-loop calibration algorithms with existing knowledge can accelerate the convergence and reduce the data volume requirements.	< 2030
System Parameter Drift	**Predictive operational stabilization**	Stabilizing system parameters, and thus stabilizing them, can help maintain high performance. The consequential reduction in the frequent need for re-calibration leads to an extension of valuable computational time.	< 2028

The synergies of these three advances will lead to the following metric expressed as **two milestones**:

1. *The operational control of a hundred-qubit processor will maintain the performance of current networks comprising up to five qubits.*

2. *The resource requirements will scale no more than polynomially in the number of qubits.*

As it is the goal to have hundreds to thousands of qubits by the year 2030, it is paramount to first reach these milestones through the means of integrative teams. We look forward to hearing from you and hope to discuss this research direction at the "ASCR Basic Research Needs in Quantum Computing and Networking" meeting with DOE program managers, the organizing committee, and the broader research community in July 2023.

[1] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos, Science **334**, 57 (2011).
[2] P. Shor, in *Proceedings of 37th Conference on Foundations of Computer Science* (1996) pp. 56–65.
[3] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, and F. G. S. L. e. a. Brandao, Nature **574**, 505 (2019).
[4] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, Science **370**, 1460 (2020).
[5] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Annual Review of Condensed Matter Physics **11**, 10.1146/annurev-conmatphys-031119-050605 (2019).
[6] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Phys. Rev. A **93**, 060302 (2016).
[7] A. Patterson, J. Rahamim, T. Tsunoda, P. Spring, S. Jebari, K. Ratter, M. Mergenthaler, G. Tancredi, B. Vlastakis, M. Esposito, and P. Leek, Phys. Rev. Applied **12**, 064013 (2019).
[8] J. Tuorila, J. Stockburger, T. Ala-Nissila, J. Ankerhold, and M. Möttönen, Phys. Rev. Research **1**, 013004 (2019).
[9] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, E. Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and J. M. Martinis, Physical Review Letters **121**, 090502 (2018).
[10] F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Phys. Rev. Applied **10**, 054062 (2018).
[11] Y. Xu, G. Huang, J. Balewski, R. K. Naik, A. Morvan, B. Mitchell, K. Nowrouzi, D. I. Santiago, and I. Siddiqi, Automatic qubit characterization and gate optimization with qubic (2021), arXiv:2104.10866 [quant-ph].
[12] F. Frank, T. Unden, J. Zoller, R. S. Said, T. Calarco, S. Montangero, B. Naydenov, and F. Jelezko, npj Quantum Information **3**, 48 (2017).
[13] P. V. Klimov, J. Kelly, J. M. Martinis, and H. Neven, The snake optimizer for learning quantum processor control parameters (2020), arXiv:2006.04594 [quant-ph].
[14] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. Carvalho, and M. J. Biercuk, PRX Quantum **2**, 040324 (2021).
[15] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, Science **303**, 1998 (2004).
[16] T. Proctor, M. Revelle, E. Nielsen, K. Rudinger, D. Lobser, P. Maunz, R. Blume-Kohout, and K. Young, Nature Communications **11**, 5396 (2020).
[17] A. Vepsäläinen, R. Winik, A. H. Karamlou, J. Braumüller, A. D. Paolo, Y. Sung, B. Kannan, M. Kjaergaard, D. K. Kim, A. J. Melville, B. M. Niedzielski, J. L. Yoder, S. Gustavsson, and W. D. Oliver, Nature Communications **13**, 1932 (2022).

The Need for Better Evaluation of Realistic Quantum Processor Performance

Meifeng Lin, Brookhaven National Laboratory, mlin@bnl.gov
Qiang Guan, Kent State University, qguan@kent.edu
Wei Xu, Brookhaven National Laboratory, xuw@bnl.gov

Topic: Codesign and integration across the quantum computing and networking stacks

Challenge In recent years, the rapid development of quantum computing technologies has led to a proliferation of diverse quantum architectures. Within the realm of digital quantum computing alone, numerous systems have emerged, including superconducting circuits, cold atoms, and trapped ions, among others. Each of these architectures has distinct strengths and limitations in relation to achieving universal digital quantum computing. Given the variability in error characteristics, qubit topology, and connectivity, assessing the quality and potential application performance of quantum systems across different architectures, or even within the same architecture, is a challenging task prior to the availability of fully fault-tolerant quantum computers **The current landscape necessitates a systematic approach, potentially through a set of performance metrics, to enable a comprehensive comparison of different quantum processing architectures.** This would allow computational scientists to evaluate the feasibility of employing targeted quantum platforms to achieve quantum advantage. Key examples of potential quantum advantage involve tackling complex computational science problems that are intractable on classical computers, such as real-time dynamics of quantum field theories, or addressing computationally demanding tasks like optimization and machine learning.

To develop a comprehensive set of performance metrics for comparing quantum architectures, several fundamental questions must be addressed. These questions are essential for ensuring the metrics accurately capture the key characteristics of quantum processors and their impact on application performance: 1). What is the relationship between the noise and error characteristics of quantum processors and the performance of applications executed on these systems? 2). Considering the dynamic nature of quantum processors, how can we assess the temporal stability of a quantum system and its ability to maintain consistent performance? 3). How can we compare different quantum processors while considering both the actual performance of applications and the range of applications that each processor can support? 4). What additional metrics will be required to quantify the performance of distributed and likely hybrid quantum systems in the future?

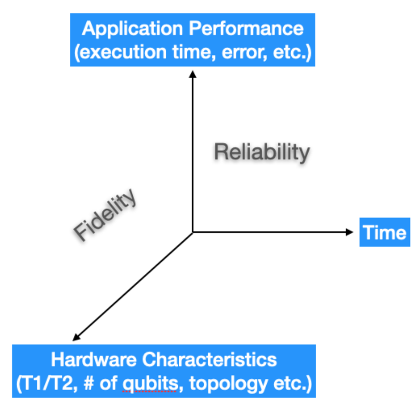

Figure 1: We emphasize the need to understand the effects of quantum hardware characteristics on application performance, as well as reliability and consistency over time.

Opportunity There is a need to develop universal metrics, potentially hierarchical, to quantify the realistic performance of diverse quantum architectures. These metrics should encompass the *fidelity, reliability* and *versatility* of quantum systems (illustrated in Figure 1), providing scientists with crucial insights into the target quantum infrastructures. In this context, *fidelity* refers to the level of errors in the application results due to imperfect qubits, *reliability* pertains to the stability of the results over time, and *versatility* reflects the types of applications a quantum system can support. Within each category, the performance metrics should further reflect granularity at the qubit, device, and system levels.

To ensure the robustness of the proposed quantum performance metrics, we need to investigate the effects of error and noise characteristics at different levels on application performance using accessible Noisy Intermediate-Scale Quantum (NISQ) platforms, such as IBM Q, Rigetti, Quantinuum, and IonQ. The potential areas of research are listed below:

1. Research to assess the effect of quantum system fidelity on application performance, including
   - **Noise characterization, modeling, and visualization** [1] of individual quantum gates and qubits, as well as the entire circuit.
   - **Quantum noise classification** to discriminate among different noise probability distributions and correlation parameters of the quantum states.
   - **Application performance prediction** on future, larger-scale quantum systems using data from current NISQ systems.
2. Research to assess the reliability of the quantum system over time, including
   - **Temporal noise prediction modeling** for leveraging exponential smoothing analysis and beyond to estimate the future noise and errors [2] and reliability over time.
   - **Bayesian analysis** to provide performance bounds based on historical NISQ data.
3. Research to assess the versatility of a given quantum system, including
   - **Uncertainty quantification and sensitivity analysis** [3] of quantum states to gauge performance differences of quantum processors and datasets.
   - **Predict performance across architectures** by leveraging the latest foundation models to gauge the architecture differences.
   - **Performance prediction for hybrid distributed quantum-classical computing**, including the communication latency, channel fidelity, and interplay between quantum and classical components and the efficiency of the algorithms used for hybrid computation.

Assessment To evaluate potential application performance, there are existing benchmarks and new ones to develop that represent promising quantum applications of interests to the Department of Energy, such as quantum chemistry and climate science. The benchmarks should also incorporate popular quantum algorithms, such as Quantum Approximate Optimization Algorithm (QAOA), Variational Quantum Eigensolvers (VQE), and others.

Timeliness/Maturity We need to have a set of well-defined and universal metrics **now** to help computational scientists evaluate and compare the utility of different quantum computers, and inform future quantum hardware design decisions, ultimately driving progress toward realizing quantum computing's full potential.

References

[1] Shaolun Ruan, Yong Wang, Weiwen Jiang, Ying Mao, and Qiang Guan. Vacsen: A visualization approach for noise awareness in quantum computing. *IEEE Transactions on Visualization and Computer Graphics*, 2022.

[2] Betis Baheri, Daniel Chen, Bo Fang, Samuel A. Stein, Vipin Chaudhary, Ying Mao, Shuai Xu, Ang Li, and Qiang Guan. Tqea: Temporal quantum error analysis. In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks - Supplemental Volume (DSN-S)*, pages 65–67, 2021.

[3] Shaolun Ruan, Ribo Yuan, Yong Wang, Yanna Lin, Ying Mao, Weiwen Jiang, Zhepeng Wang, Wei Xu, and Qiang Guan. VENUS: A Geometrical Representation for Quantum State Visualization. 3 2023.

Enhancing noise-awareness in quantum compilers

Ji Liu, Zain Saleem, Paul Hovland
Argonne National Laboratory
(ji.liu@anl.gov)

Topic: Quantum Compilation

Challenge: Quantum compilers play a critical role in practical quantum compilation, particularly in the noisy intermediate-scale quantum era. With a limited number of qubits/connectivity and the presence of noisy quantum operations, it is crucial to leverage the quantum compiler to optimize circuits and map them onto the target hardware.

State-of-the-art compilers often focus on reducing the number of quantum gates due to its direct association with circuit time and noise. The impact of noise on quantum compilation remains an active research area, and the overall effectiveness of noise-mitigation strategies still needs to be fully understood. Some studies concentrate exclusively on reducing noise at specific stages of the compilation process or focus on one specific type of error source. Noise-aware compilations rely on relatively simple noise models, such as the estimated-success-probability(ESP) noise model [NPS+20] employed by noise-adaptive mapping and routing algorithms, which is known to be inaccurate. The short compilation time requirement also limits the complexity of the noise model employed at runtime.

Incorporating noise information for pulse-level engineering presents another significant challenge. While pulse-level optimization has gained significant attention due to its potential to reduce circuit time, it faces a fundamental problem: the system Hamiltonian is hard to model because of the presence of interaction terms such as cross-talk. Despite several gradient-based pulse-level methods being proposed for parameterized circuits, the complex nature of quantum control means that these have yet to be demonstrated on hardware with real-world noise.

Another significant challenge is incorporating error mitigation techniques in the compilation flow. Various approaches have been proposed to mitigate different types of errors, such as dynamical decoupling methods for idling error mitigation, randomized compiling for converting coherent errors into stochastic noise, and ensemble approaches for mitigating machine-dependent noise. There is still a limited understanding of the overall and collective effectiveness of these approaches, and a systematic view of their integration is lacking. Additionally, incorporating error mitigation techniques into the compilation flow can introduce both compilation and circuit execution overhead. Effectively combining these error mitigation techniques to collectively reduce noise while minimizing compilation and execution overhead is therefore important.

Opportunity: An efficient real-time noise estimator is essential for building a noise-aware compiler. Unlike noise simulators that simulate noise in the system, a noise estimator predicts the noise impact on circuit fidelity given an input circuit and a target system. It can significantly improve the speed of noise prediction compared to noise simulators, making it well-suited for quantum compilers that evaluate different circuit designs at runtime. To this end, several ideas have been proposed based on neural network and graph transformer models [LZ20, WLC+22].

Noise simulators can still be valuable for designing and testing noise-aware compilers, guiding gate decomposition, and optimizations. In addition to gate-level noise simulation, it is crucial for investigating pulse-level noise simulation, which has received less attention despite the increasing interest in control-pulse-level optimization. While there are many software packages (e.g., Qiskit, Cirq, PyQuil, PennyLane) that provide noise simulation, the noise is simulated at the gate level by probabilistically inserting random Pauli gates or Kraus Operators to describe a noisy quantum channel. However, simulating noise at the level of time evolution allows for a more realistic study of quantum circuits with noise based on the physical model [LAS+22]. Currently, there needs to be more integration with real device calibration data, and further research is needed to determine how pulse-level noise simulation can guide compiler design.

Even after having a good noise model, system noise drifting may significantly damage the accuracy of noise models. Quantum systems are periodically calibrated, and the calibration data is typically used to guide the real-time noise estimator and noise-aware compilation. Significant drifting may occur within the same calibration cycle, rendering the calibrated data unreliable. It has been demonstrated that error calibration data collected just before circuit execution can significantly improve the results of noise-aware compilation [WSM20]. However, collecting such calibration data incurs a noticeable overhead. Unstable system noise can also adversely affect error mitigation techniques such as zero noise extrapolation, which requires stable noise as circuit folding factors increase. Collaborative efforts between computer scientists and device experts are critical for effective hardware and software co-design. For example, the recently introduced ibm_sherbrooke device leverages a new calibration strategy that prioritizes more stable system performance with minimal recalibration needs instead of minimizing error rates at the expense of stability, which makes it suitable for noise-aware compilers and error mitigation techniques.

The challenge of incorporating error mitigation techniques can be tackled by systematically evaluating different noise mitigation approaches and studying their impact throughout the entire compilation flow, rather than focusing on individual stages. Future research should prioritize the integration of various error mitigation techniques and collective approaches that minimize the overall compilation and execution overhead.

Assessment: Standard circuit fidelity metrics are not enough to evaluate potential solutions. As the noise-aware compilation and error mitigation methods can introduce significant overhead, it is crucial also to consider the total quantum and classical execution time. Therefore, metrics like CLOPS [WPJA+21] are essential. By evaluating noise-aware compilation methods with a combination of fidelity and execution time metrics, a comprehensive analysis can be performed to determine the most effective solution.

Timeliness or maturity: Hardware advances are essential for the development of quantum computing, but the full potential of these systems can only be realized through effective software solutions. While noise-aware compilation techniques have been developed, there is a need for improving the accuracy of the noise models and also model noise at the control-pulse level. In particular, there has been considerable progress in developing publicly available software tools for qubit control, leading to several recent advancements in optimizing the system at the control pulse level. However there is still much work to be done to study noise-aware pulse compilation in real hardware systems. Additionally, while there have been significant recent advancements in error mitigation techniques, more systematic studies of how to integrate these approaches into the software stack need to be done. It is thus an opportune time to focus on systematic noise-aware compilation research, which could help unleash the full potential of quantum hardware while fully utilizing error mitigation techniques.

References

[LAS+22] Boxi Li, Shahnawaz Ahmed, Sidhant Saraogi, Neill Lambert, Franco Nori, Alexander Pitchford, and Nathan Shammah. Pulse-level noisy quantum circuits with qutip. *Quantum*, 6:630, 2022.

[LZ20] Ji Liu and Huiyang Zhou. Reliability modeling of nisq-era quantum computers. In *2020 IEEE international symposium on workload characterization (IISWC)*, pages 94–105. IEEE, 2020.

[WLC+22] Hanrui Wang, Pengyu Liu, Jinglei Cheng, Zhiding Liang, Jiaqi Gu, Zirui Li, Yongshan Ding, Weiwen Jiang, Yiyu Shi, Xuehai Qian, et al. Quest: Graph transformer for quantum circuit reliability estimation. *arXiv preprint arXiv:2210.16724*, 2022.

[WPJA+21] Andrew Wack, Hanhee Paik, Ali Javadi-Abhari, Petar Jurcevic, Ismael Faro, Jay M Gambetta, and Blake R Johnson. Quality, speed, and scale: three key attributes to measure the performance of near-term quantum computers. *arXiv preprint arXiv:2110.14108*, 2021.

[WSM20] Ellis Wilson, Sudhakar Singh, and Frank Mueller. Just-in-time quantum circuit transpilation reduces noise. In *2020 IEEE international conference on quantum computing and engineering (QCE)*, pages 345–355. IEEE, 2020.

Quantum algorithms for training large-scale machine-learning models

PI: Junyu Liu, Postdoctoral Fellow, University of Chicago, junyuliu@uchicago.edu
Co-PI: Yuri Alexeev, Principal Project Specialist, Argonne National Laboratory
Co-PI: Liang Jiang, Professor, University of Chicago

Call: Position Paper for Workshop on Quantum Computing and Networking
Topic: Algorithms

Large-scale machine learning models, such as GPT-3[1,2], hold immense promise for societal advancement, given their capacity to handle and scrutinize enormous volumes of data efficiently. Leveraging sophisticated algorithms and remarkable computational prowess, these machine learning models can reveal previously elusive insights and patterns. The implications are profound, potentially catalyzing breakthroughs and enhancing results across various sectors such as healthcare, transportation, and finance, among others.

Yet, the extensive parameter training of these models incurs significant costs and contributes to substantial carbon emissions. To exemplify, the training of GPT-3 led to an expenditure of twelve million dollars and generated over five-hundred tons of CO_2 equivalent emissions[3]. Therefore, it is crucial to strive for sustainability and efficiency in large-scale machine learning models, including large language models (LLMs). As the future is anticipated to see the training of numerous specialized GPT models, the energy consumption and resultant harmful atmospheric emissions from computational centers are set to escalate dramatically.

Opportunity: Future training of large-scale machine learning models could potentially be addressed by employing quantum machine learning[4] on quantum computers. This approach, where machine learning algorithms are executed on quantum devices, is broadly perceived as a highly advantageous application of quantum algorithms. However, despite accelerated development and noteworthy progress, existing quantum machine learning algorithms are encumbered with significant limitations in both theory and practice. Practical applications of these algorithms for near-term devices often lack the theoretical foundation that could guarantee or convincingly suggest a superior performance over their classical equivalents. Moreover, in the context of fault-tolerant scenarios in quantum machine learning[5–13], it is indeed possible to demonstrate rigorous super-polynomial quantum speedups for highly structured problems[14–16]. Nonetheless, these propositions are arguably still a considerable distance from the practical, state-of-the-art applications of classical machine learning. Some of these methods primarily utilize quantum states as training data as opposed to conventional data[15,17–20]. While these approaches are promising, they arguably do not align with the most critical current applications of classical machine learning. Therefore, it is essential to expand our comprehension of quantum machine learning. This involves understanding the theoretical assurances it could potentially provide and the practical, timely problems of classical machine learning it could resolve, at least in principle. For example, it should be pertinent to scalable and sustainable challenges inherent in large-scale machine learning."

Assessment: We aim to address the previously discussed challenges by developing comprehensive quantum machine learning algorithms. These algorithms are designed to be relevant to the current machine learning community and are somewhat fortified with theoretical assurances. We observe that once a significant portion of neural network training parameters has undergone sparse training and the classical training parameters have been compiled to a quantum computer, a quantum enhancement can be identified early in the training phase, prior to the exponential growth of error[21–24]. The crux of our approach involves adapting the quantum algorithm, which solves differential equations[25], to execute (stochastic) gradient descent algorithms – arguably the principal classical machine learning algorithm – on a quantum processor post-linearization. The anticipation of a potential quantum enhancement is based on the implementation of a variant of the *Harrow-Hassidim-Lloyd* (HHL) algorithm[26], an efficient quantum technique for sparse matrix inversion that solves the problem within a logarithmic time scale for suitably conditioned sparse matrices. We discovered that[27] our algorithm can solve machine learning problems of large-scale model-dimension

n within a polylogarithmic time scale multiplied by the number of iterations T, or the square of T. This scaling in n surpasses any known classical algorithms. However, while it outperforms in this area, there is no assurance that our hybrid quantum-classical algorithm will invariably outperform all other potential classical algorithms for related, albeit different tasks, such as non-gradient-based algorithms. As a result, to the best of our understanding, our research indicates the possibility of a substantial quantum speedup or enhancement of specific classical algorithms, rather than a quantum advantage over the entire problem class.

Timeliness or maturity: Our proposed research trajectory could considerably enhance the scalability and sustainability of classical large-scale machine learning models. We support this assertion with numerical evidence up to 103 million training parameters. This work offers robust theoretical assurances and intersects with cutting-edge classical machine learning research. Our approach diverges significantly from the thinking behind variational quantum algorithms. Instead, it seeks to bolster classical machine learning with a crucial quantum step that forms a bottleneck for classical training. In essence, it can be perceived as substantiating the expectation that quantum versions of neural networks could give rise to novel computational tools[28].

Our research is anticipated to pave the way for numerous potential advancements in quantum machine learning, where there is a realistic expectation for algorithmic enhancements. In our supplemental material, we highlight several potential research paths that could yield particularly beneficial results. These include developing a time-dependent version during gradient descent trajectories, establishing improved formal criteria for dissipation, exploring connections to diffusion models in classical machine learning and LLMs[29], enhancing the truncated HHL algorithms theoretically, and identifying mechanisms of possible quantum speedups beyond the concept of dissipation. It is our hope that our work can inspire further research in these areas.

References:

[1] T. Brown et al. *Adv. Neur. Inf. Process Sys.* 33, 2020. arXiv:2005.14165. [2] K. Roose. *The New York Times*, 2022. [3] D. Patterson et al. arXiv:2104.10350. 2021. [4] J. Biamonte et al. *Nature* 549, 2017. arXiv:1611.09347. [5] A. Peruzzo et al. *Nature Comm.* 5, 2014. arXiv:1304.3061. [6] J. R. McClean et al. *New J. Phys.* 18, 2016. arXiv:1509.04279. [7] S. McArdle et al. *Rev. Mod. Phys.* 92, 2020. arXiv:1808.10402. [8] M. Cerezo et al. *Nature Rev. Phys.* 2021. arXiv:2012.09265. [9] E. Farhi et al. arXiv:1411.4028. 2014. [10] S. Ebadi et al. *Science* 376, 2022. arXiv:2202.09372. [11] E. Farhi and H. Neven. arXiv:1802.06002. 2018. [12] V. Havl. *Nature* 567, 2019. arXiv:1804.11326. [13] J. Liu et al. arXiv:2210.06723. 2022. [14] R. Sweke et al. *Quantum* 5, 2021. arXiv:2007.14451. [15] Y. Liu et al. *Nature Phys.* 17, 2021. arXiv:2010.02174. [16] N. Pirnay et al. *Phys. Rev. A* 107, 2023. arXiv:2210.14936. [17] S. Lloyd et al. *Nature Phys.* 10, 2014. arXiv:1307.0401. [18] H.-Y. Huang et al. *Nature Comm.* 12, 2021. arXiv:2011.01938. [19] H.-Y. Huang et al. *Nature Phys.* 16, 2020. arXiv:2002.08953. [20] H.-Y. Huang et al. *Science* 376, 2022. arXiv:2112.00778. [21] T. Hoefler et al. *J. Mach. Learn. Res.* 22, 2021. arXiv:2102.00554. [22] N. Lee et al. *International Conference on Learning Representations.* 2019. [23] C. Wang et al. *International Conference on Learning Representations.* 2020. [24] H. Tanaka et al. *Adv. Neur. Inf. Process Sys.* 33, 2020. arXiv:2006.05467. [25] J.-P. Liu et al. *Proc. Natl. Ac. Sc.* 118, 2021. arXiv:2011.03185. [26] A. W. Harrow et al. *Phys. Rev. Lett.* 103, 2009. arXiv:0811.3171. [27] J. Liu et al., Mar. 2023. arXiv: `2303.03428 [quant-ph]`. [28] R. Bondesan and M. Welling, 2021. arXiv:2103.04913. [29] J. Ho et al. *Adv. Neur. Inf. Proc. Sys.* 33, 2020. arXiv:2006.11239.

Quantum Networking at the Mesoscale: From Modular Quantum Processors to the Quantum Datacenter

Authors: Michael J. Martin (mmartin@lanl.gov)

Topic: Codesign and Integration

Challenge: Neutral atom quantum computers use large arrays of controlled, individually-trapped neutral atoms as qubits, with short-range interactions for entanglement provided by excitation to Rydberg states. Already, making optimal use of the hundreds of atoms in state-of-the-art neutral atom systems requires networking at the microscale. For example, recent work [1] has shown that entanglement transport based on moving individual atoms or sub-arrays of neutral atoms is a powerful resource for generating cluster states and implementing quantum error correcting codes. Because coherence is preserved under transport, networking via these micron-scale, physical-space moves is perhaps the most promising way to advance the state of the art in these processors towards fault tolerant quantum computing, powerful quantum simulators, and adiabatic quantum computing. However, scaling these advances towards systems beyond the few-hundred qubit scale requires new networking approaches to realize modular, scalable architectures of networked neutral atom processors—a network scale we call the mesoscale to distinguish it from long-distance quantum networks. Other applications include hybrid systems of networked sensors/processors and analog simulators/processors [2]. With meso-scale networks established for "quantum data centers" at the meter-scale, these systems will be further poised to integrate into long-distance quantum networks for distributed, high-performance quantum computing and sensing as quantum transduction schemes become available.

Opportunity: Here, we envision an architecture where arrays of up to several hundred atoms are networked to nearby arrays of comparable size (Fig. 1). The number of atoms in each sub-array (i.e., a single quantum processor) is set by the atomic spacing, which is typically limited to the few micron scale; the field of view of the objective lens used for trapping; and the optical power to each objective lens. Presently, state-of-the-art systems comprise a two-dimensional array, trapped and imaged by a single objective lens, with array sizes limited to the several hundreds of atoms. This is a limit to scaling to larger systems of qubits, which will be necessary to realize the benefit of fault-tolerant quantum computing for scientific research, which prevents scaling to very large systems of qubits. The opportunity proposed here is to link multiple sub-arrays together, each with its own objective lens and laser system, which are required for trapping and manipulating atoms, in a modular architecture.

Success would open the door to (1) novel quantum sensing approaches, where variationally-optimized probe states in a large, distributed system would yield quantum sensing advantage below the standard quantum limit; (2) hybrid processor/memory architectures, where different atomic species could be used to process and store quantum information, respectively; and (3) a modular approach to scaling neutral atom quantum computers beyond the few-hundred qubit scale, whose modular design will integrate well into long-distance quantum networks with the appropriate quantum transduction schemes.

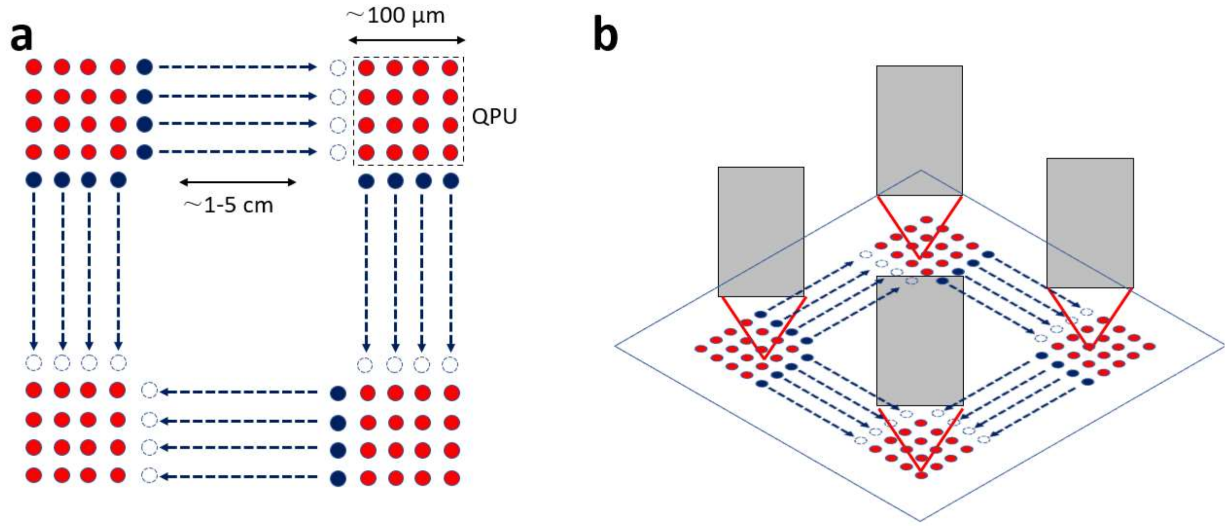

Figure 1 (a) Proposed methodology for connecting neutral atom-based quantum processing units (QPUs, red circles) via messenger atoms (blue circles) that are shuttled between processors over cm-scale distances. While each quantum processor would contain hundreds of atoms, we illustrate them with far fewer atoms here for clarity. (b) Architecture for scaling to multi-processor system with external imaging optics (grey). The separation between processors allows space for external optics that are required for trapping and imaging atoms within each processor.

Assessment: Success in this area could be assessed incrementally. First milestones would be the demonstration of deterministic atom transport between small quantum processors that each have their own optics, thus demonstrating the promise of the modular approach. The actual methodology for transport is a rich research question, with numerous open research questions that are ripe for exploration. Follow-on steps would be demonstration of qubit coherence, and demonstration of remote entanglement in few-qubit systems.

Timeliness or maturity: There has been a tremendous amount of industry and academic interest and success in developing neutral atom-based quantum processors. These systems have the natural advantage of long qubit coherence times, and relatively easy scaling to the few-hundred qubit scale at the single-quantum processor level. In this forward looking Position Paper, we envision the next steps that will be necessary to fully realize the quantum advantage of these systems, which will require quantum networking beyond the microscale, to the mesoscale and beyond. The impact of the proposed research path would be to unlock the potential of neutral atom quantum computing toward universal quantum computing.

References

[1] D. Bluvstein and et al., "A quantum processor based on coherent transport of entangled atom arrays," *Nature,* vol. 604, pp. 451-456, 2022.

[2] S. Ebadi and et al., "Quantum optimization of maximum independent set using Rydberg atom arrays," *Science,* vol. 376, no. 6598, pp. 1209-1215, 2022.

Space-based quantum networks are an essential component of future architecture for distributed quantum computers and quantum-enhanced secure communication

Kristina Meier,[1*] Raymond Newell,[2] Nicholas Dallmann[2]

[1] *Intelligence and Space Research Division, Los Alamos National Laboratory*
[2] *Materials Physics Applications - Quantum Division, Los Alamos National Laboratory*
[*] *corresponding author: kameier@lanl.gov*

Abstract: A successful US quantum satellite would require large investment, national priority, and a diverse set of expertise. But, it would deliver US-owned quantum links that would allow for quantum-enhanced secure communications and the ability to connect quantum computers over long distances.

1. Topic

Models, Codesign and Integration, and Applications

2. Challenge

Sending information between quantum computers, or linking them for distributed quantum computing, requires sending quantum information over long distances. It would represent a significant step forward in quantum information technology if the US had the capability to connect quantum computers on the east coast, e.g., IBM, with ones on the west coast, e.g., Google. Furthermore, part of the ASCR mission is to "...develop, and deploy computational and networking capabilities...," and as the world moves towards practical quantum technology, quantum computers fit squarely into ASCR's mission. The need for long-distance (space-based) quantum links is, as such, an indispensable part of realizing this mission.

Non-quantum communication is done with telecom fiber networks that implement amplifiers and repeaters for copying and/or boosting the classical signal to mitigate fiber losses. Unfortunately for quantum signals, the signals are ultra-low power (single-photon level) and cannot be copied due to the no-cloning theorem [1]. This means that it would take tens of billions of years to be successful in transmitting one physical bit of quantum information (qubit) through fiber from the eastern to the western sides of Texas given a standard fiber loss of 0.2 dB/km.

Implementing free-space networks using satellites would allow for long-distance transmission of single photons at more practical rates. J.Y. *et al.* (2017) showed that a satellite-based system can operate with losses of 64-82 dB with then-current technology over a quantum link equivalent to 1200 km [2]. Compare this with the 240 dB loss within a fiber for the equivalent link length, and satellite-based quantum links offer an 10^{17} times improvement, enabling high-value functions, such as Quantum Key Distribution. Furthermore, one could imagine a constellation of satellites that can transmit to either downlinks or to other satellites within the constellation, depending on the destination of the information. This would allow for the photons to do most of their travel via low-atmosphere links, instead of lossy fiber.

Long-range distribution of quantum entanglement, continental-scale or beyond, would present unique opportunities for scientific discovery. In addition to the transformative power of networked quantum computers, which may be expected to obey scaling at least as fast as Metcalfe's Law [3], applications in quantum sensing can be explored. Astronomy with distributed aperture optical telescopes has been proposed [4, 5]. Furthermore, experimental investigation of quantum gravity may help resolve long-standing conflicts between quantum field theory and general relativity [6]. This research would require access to spaceborne quantum sources and receivers.

3. Opportunity

Satellite-based quantum networks would allow quantum information transfer over extremely long distances and, with improved technology, could reach practical rates for communication between quantum computers. This work must begin with a quantum space mission to demonstrate practical space-based quantum links and must be a national priority. The following subsections illustrate steps to a successful quantum space mission that would begin to crack the challenge of low-loss quantum links between quantum computers.

(1) Model: The first step in a quantum space mission would be to develop link budgets to model loss within quantum links for multiple architectures, e.g., ground-to-satellite (uplink), satellite-to-ground (downlink), satellite-to-satellite (crosslink), using photon rates reasonable for current photon source technology, i.e., spontaneous parametric [7] or quantum nano-emitters [8]. This model would also quantify the required link rates to support connections between quantum computers. This modeling would give insight into the optimal architecture to implement for a space-based quantum network.

(2) Ground-based validation of model with lab-like hardware: Once an architecture is chosen, the optimal photonic sources and detectors are developed. These choices are based on the brightness of the source, the detection efficiency of the detector, the required size, weight, and power (SWaP) of supporting the source or detector either in space or on the ground, and how stable the source is, e.g., how many moving parts are there? Can the source be integrated? Does the source require active alignment onboard the satellite? Atmospheric turbulence and loss should also be recreated in the lab or on optical test ranges to test qubit decoherence over free-space links.

(3) Validation with flight-like hardware: Advanced knowledge about space-ready components is essential for this step. What parts of the lab-based model need to be replaced with space-ready components? Do our experiments still work in the lab with flight hardware? Can a proof-of-principle flight be accomplished on a CubeSat?

(4) Space qualification of flight-like hardware: Are there methods for integrating photonic sources and detectors that would allow for a quantum satellite that is robust to vibration testing? Is there ample information on the effects of space radiation on the materials that make up both photonic sources and detectors? With current technology, how long would a quantum satellite last in orbit before needing to be replaced due to radiation damage?

(5) Flight: The final step is to develop a team for monitoring and detecting photons from the satellite and/or sending photons to a detector on the satellite. This would require expertise in pointing and tracking for faint signals from satellites.

For successful completion of a quantum space mission, collaboration between groups with diverse expertise is required. This could be internal collaborations for organizations with diverse portfolios, and/or external collaborations. An ideal team would include experts in: quantum information science; quantum sources; photonic detectors; space-readiness; pointing, tracking, and acquisition; nanotechnology; electrical engineering; satellite engineering; and systems engineering.

4. Assessment

A successful quantum space mission would be one where physical qubits are transferred over a quantum link with high fidelity, i.e., minimal information loss. High fidelity is achieved with high-efficiency photonic sources, optimal pointing and tracking, sufficiently large telescope apertures for ground detection, and high-efficiency detectors. Bringing this capability to the United States is paramount.

5. Timeliness or Maturity

A US-based quantum satellite mission is long overdue. In 2017, China launched a quantum satellite (Micius) and demonstrated both entanglement distribution and quantum key distribution [2]. Canada is working on their Quantum Encryption and Science Satellite (QEYSSat) with an expected launch in 2024-2025 [9]. One application of their satellite is to create a transcontinental quantum link between Canada and Europe [10]. In the United States, there is currently no large-scale quantum satellite mission, hindering the United States' lead in quantum technologies and creating a potential future national security concern if other countries greatly surpass the nation's abilities for sending secure information between quantum computers. Given advances in nanoemitter-based quantum sources [8] and CubeSats driving the low-cost space revolution, now is the time to form a US-based quantum satellite mission.

References

1. W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," Nature **299**, 802–803 (1982).
2. J. Yin et al., "Satellite-based entanglement distribution over 1200 kilometers," Science **356**, 1140–1144 (2017).
3. B. Metcalfe, "Metcalfe's law after 40 years of ethernet," Computer **46**, 26–31 (2013).
4. D. Gottesman, T. Jennewein, and S. Croke, "Longer-baseline telescopes using quantum repeaters," Physical Review Letters **109** (2011).
5. J. Crawford et al., "Towards Quantum Telescopes: Demonstration of a Two-Photon Interferometer for Quantum-Assisted Astronomy," arXiv e-prints arXiv:2301.07042 (2023).
6. M. Mohageg et al., "The deep space quantum link: prospective fundamental physics experiments using long-baseline quantum optics," EPJ Quantum Technology **9**, 25 (2022).
7. M. Bock, A. Lenhard, C. Chunnilall, and C. Becher, "Highly efficient heralded single-photon source for telecom wavelengths based on a ppln waveguide," Opt. Express **24**, 23992–24001 (2016).
8. N. Tomm et al., "A bright and fast source of coherent single photons," Nature Nanotechnology **16**, 399–403 (2021).
9. "Quantum encryption and science satellite (qeyssat)," Canadian Space Agency (2023).
10. "Connecting canada and europe through quantum satellite communication," Waterloo News (2023).

Quantum Co-design and Integration for Science Applications

S. M. Mniszewski
Los Alamos National Laboratory
smm@lanl.gov

Topic: This paper addresses co-design and integration of algorithms across applications.

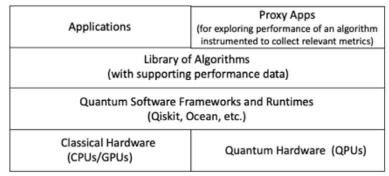

Fig.1. SW/HW stack including Proxy Apps and library.

Challenge: As quantum computing becomes part of the HPC landscape [1], applications that are relevant to DOE will adopt quantum acceleration with the promise of quantum speedup over purely classical implementations. Due to the limited number of qubits, connectivity, and short coherence times hybrid quantum-classical approaches are necessary for accelerating applications today [2-3]. For gate-based quantum computing architectures such as IBM Q, Ion Q, and others, the variational quantum eigensolver (VQE) and the quantum approximate optimization algorithm (QAOA) can be used. For quantum annealers, such as D-Wave, *qbsolv* and user/vendor hybrid solvers can be used. Classical cluster-based CPU/GPU codes can also be used in both cases. The component algorithms that make up an entire application will have quantum or hybrid form. They are typically user-based and require an ansatz (gate-based) or embedding (quantum annealer) and error mitigation when run. Some examples include solving linear systems, quantum Fourier transform, electronic structure, graph decomposition, and more. There can be many different hybrid implementations depending on the problem to be run and the architecture to be run on. Each algorithm represents a family of implementations defined by a set of problem, architecture, and performance-relevant parameters. This does not exist today. We have limited understanding of how performant the algorithms we develop are and how they would work in extended situations.

Opportunity: Quantum co-design approaches [4-5] come in many forms. Drawing from the experience from the Exascale Computing Project (ECP) [6] co-design capabilities would focus on libraries and proxy applications (or proxy apps) to identify, collect, and develop common hybrid quantum-classical algorithms for different application categories such as chemistry, materials, machine learning, etc. A proxy app for each algorithm or a set of algorithms would be used to explore implementation options and characterize and determine performant implementation(s) on quantum architectures. Where this fits in the quantum SW/HW stack is shown in Fig. 1. These algorithms would require a common API and an evolving form for current NISQ architectures and may even help identify improvements in the quantum software stack and quantum hardware. These can also serve as useful benchmarks for testing new quantum architectures. This would require quantum algorithm developers, software stack developers, and hardware designers working together to determine the best solutions.

Assessment: Success is measured by many factors such as time to solution, efficient use of error mitigation, circuit width and depth, estimates of quantum resources required, and scalability. If it can be made easier to develop applications that use quantum or hybrid components that is a win. These libraries and supporting proxy apps would help determine the best component algorithms to use for an application on a chosen NISQ architecture. It could potentially show that a combination of different algorithm-architecture pairs could provide even better performance.

Timeliness or maturity: Due to the many algorithms that have been developed, tested, and documented through publications and github repositories, the time is right to start collecting this knowledge in the form of libraries and proxy apps. Having libraries of algorithms that are in a sense *certified* per quantum architecture for performance, resources required, error resilience, sensitivity to noise, and scalability would allow development of new quantum applications and new algorithms easier. Note that these algorithms would need to be parameterized based on problem size range and other architecture-based relevant information.

Where would such an effort begin? Teams working in common application areas can bring their component algorithms together for performance characterization. These algorithms would start as hybrid quantum-classical versions of familiar classical algorithms. As new ideas for quantum algorithms emerge based on quantum principles that were not possible classically, they will also be included. Proxy Apps and other software tools would be needed to run and collect timing and efficiency results. These results may be distributions. The collection of component algorithms would become a library with supporting characterization information. One can envision these libraries of quantum algorithms becoming standard and being supported by teams of quantum computer scientists that continually keep them current.

As quantum architectures evolve towards post-NISQ with larger sizes, increased connectivity, longer coherence times, and error mitigation/correction the library algorithms would also need to evolve. Having such a co-design capability would allow for building a new application code with component quantum and/or hybrid algorithms with known performance. This approach will make it possible to estimate the potential for quantum speedup.

References

1. National Academies of Sciences, Engineering, and Medicine, 2023, Charting a Path in a Shifting Technical and Geopolitical Landscape: Post-Exascale Computing for the National Nuclear Security Administration, Washington, DC: The National Academies Press. https://doi.org/10.17226/26916.
2. R. Shaydulin, H. Ushijima, C. F. A. Negre, I. Safro, S. M. Mniszewski, Y. Alexeev, 2019, A Hybrid Approach for Solving Optimization Problems on Small Quantum Computers, *IEEE Computer,* 52 (6), 18-26. https://doi.org/10.1109/MC.2019.2908942
3. J. E. Moussa, A. D. Baczewski, M. Sarovar, 2017, Co-design Strategies for Quantum Simulators, *Sandia Report* SAND2017-0164C.
4. G. Li, A. Wu, Y. Shi, A. Javadi-Abhari, Y. Ding, Y. Xie, 2021, On the Co-Design of Quantum Software and Hardware, 8th Annual ACM International Conference on Nanoscale Computing and Communication (NANOCOM ’21). ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3477206.3477464
5. T. Tomesh, M. Martonosi, 2021, Quantum Codesign, *IEEE Micro*, 33-40. https://doi.org/10.1109/MM.2021.3094461
6. Alexander F., et al., 2020, Exascale Applications: Skin in the Game. *Phil. Trans. R. Soc. A* 378: 20190056. http://dx.doi.org/10.1098/rsta.2019.0056

Benchmarking early fault-tolerant quantum computers

April 2023 **Topic: Error correction and mitigation**

Alexandra E. Moylett and Joan Camps, Riverlane ⟨{alex.moylett, joan.camps}@riverlane.com⟩

Challenge The advent of the era of early fault-tolerant quantum (EFTQ) computers enables, for the first time, the trade-off between qubit count and fidelity. At the fidelity allowed in the era of noisy intermediate-scale quantum computation (NISQ), the maximum circuit depth that can be run is of a few tens of gates. Despite significant efforts deployed in the algorithmic front, and the development of statistical methods to mitigate errors [1], this maximum depth has limited the usefulness of NISQ computers. The NISQ era has seen a convergence in the benchmarking of quantum computers. The reporting of the performance of NISQ quantum computers has become standarized in terms of qubit coherence times, gate fidelities, and their tomographic properties [2].

In parallel with NISQ algorithms, the last decade has seen steady development and optimization of fault-tolerant algorithms – bringing down the cost of running applications of interest by orders of magnitude [3]. Large-scale applications such as drug discovery and factoring 2048-bit RSA integers will still require resources beyond what we have in EFTQ [3]–[5]. These developments typically assume an arbitrarily high qubits-to-fidelity trade-off, ie as many qubits available as necessary to prevent a logical error.

In contrast, the optimization and development of algorithms that may run on EFTQ computers, where we have limited fidelity due to a limited physical qubit count, remains largely unexplored. In particular, EFTQ are expected to have access to high fidelity Clifford operations but a very limited supply of noisy non-Clifford operations. The seeds of ETFQ exploration appear in recent IBM work on the interplay between error correction and error mitigation techniques [6] and statistical quantum algorithms [7], [8]. Yet there remain many unanswered questions: how best should be benchmark EFTQ, what are the first experiments we should run, which might show quantum advantage and for which is this impossible.

On EFTQ benchmarking, initial metrics for EFTQ performance include the logical error rate of error-corrected qubits [9], which is a good proxy for logical coherence times. One might expect that next we need to perform logical gates and measure their fidelity. However, the logical operation of fault-tolerant quantum computers can often be thought of consisting of elementary operations that are smaller than gates – for instance, when operating the surface code with lattice surgery [10], logical patch motion, merging, and splitting are all part of logical gates. While there are some proposals to benchmark some such logical primitives of fault-tolerant quantum computers [11], there is no consensus regarding what these primitives should be and how to report their performance.

Opportunity Early fault-tolerant quantum computers open opportunities for exploring small-scale applications. There is no consensus, and few suggestions, for what to do with these early devices. A window of opportunity presents itself to explore and define a roadmap of EFTQ experiments. Error correction and error mitigation techniques can lead to proof-of-principle applications and toy demonstrations we can implement on a small fault-tolerant machine [6], thus making the most of the qubits-to-fidelity trade-off. While some seeds have been planted, the full arsenal of relevant tools have not yet been applied to EFTQ. Such tools include error mitigation theory, magic resource theory (including convex optimisation tools) [12] and random compilation theory [13].

Utilizing these techniques will require a detailed understanding of the noise that affects both physical and logical qubits. This requires the development of benchmarks to learn not just the probability of a logical error occurring, but also the types of logical errors that occur in practice [14]. Such benchmarks will initially be applied to quantum memory experiments like the ones that exist today [9], [15], but can also be used to study how noise affects key algorithmic components. For example, gate compilation techniques often produce common sequences of logical gates such as alternating sequences between Hadamard and phase gates [16], which can benefit from quantum benchmarking. Another example of where benchmarks can be utilized is primitives in lattice surgery [10].

Finally, any developments must be made in collaboration with quantum hardware companies, to optimize for the native gates, connectivity, and physical noise profile of their processors as well as understand directly the performance of error correction primitives on their hardware. This includes developments across the whole quantum computing stack, from the design of logical circuits, to benchmarks, quantum error correction and mitigation techniques, and decoding algorithms [17].

Assessment The success of this program may be measured in terms of adoption and attainment of newly defined EFTQ milestones, the development of new optimisations across the stack targeting improved benchmarks (including, *e.g.*, improved gate fidelity, decoder optimizations), acceleration of roadmaps towards such milestones, and convergence of practice towards meaningful metrics for algorithmic performance in EFTQ devices. An early goal for success is the widespread adoption for benchmarks specific to the EFTQ regime, which will be superior to NISQ benchmarks so long as *(i)* they inform accurate predictions for the success of computations, *(ii)* their improvement unlocks more valuable computations, and *(iii)* they are used to optimise algorithms for the machines they run on. We expect the shift from EFTQ to full fault tolerance to be gradual, therefore it is important to extract the maximum benefit from the EFTQ era. Major successes in this regard would include finding applications of EFTQ beyond proof-of-principle demonstrations. Perhaps this is more likely in academic toy problems rather than industry applications. But even toy demonstrations can capture the imagination of the public and leap the whole field forwards.

Timeliness or maturity Quantum computers have reached a stage of maturity which has transformed fault tolerance from a theoretical concept to experimentally demonstrable reality. In the last few years we have seen significant breakthroughs in experimental quantum error correction, from implementation of logical operations to a demonstration of suppressing errors using an error-correcting code [9], [15], [18]–[20]. At the same time, algorithmic improvements and resource estimation have given us rigorous upper bounds on the resources required for large-scale fault-tolerant quantum computers to solve significant problems of interest [3]–[5]. Multiple hardware manufacturers have now developed roadmaps for developing quantum processors at scale [21]–[24]. As we begin to see the rise of small-scale quantum computers with modest error correction, it is important for us to develop a roadmap for the rest of the quantum computing stack to reach the long-term goal of large-scale fault tolerance. Through the development of benchmarks and mitigation techniques for the EFTQ era, we will understand both the potential and the key bottlenecks in real-world quantum computers as the number of high-fidelity logical qubits increases. From this, we can lay forth the key steps and milestones which will guide us to practical large-scale quantum computation.

References

[1] Z. Cai, R. Babbush, S. C. Benjamin, *et al.*, *Quantum error mitigation*, 2022. arXiv: 2210.00921 [quant-ph].

[2] K. Wright, K. M. Beck, S. Debnath, *et al.*, "Benchmarking an 11-qubit quantum computer," *Nature Communications*, vol. 10, no. 1, p. 5464, 2019.

[3] J. Lee, D. W. Berry, C. Gidney, *et al.*, "Even more efficient quantum computations of chemistry through tensor hypercontraction," en, *PRX Quantum*, vol. 2, no. 3, p. 030 305, 2021, arXiv: 2011.03494.

[4] C. Gidney and M. Ekerå, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits," *Quantum*, vol. 5, p. 433, 2021.

[5] N. S. Blunt, J. Camps, O. Crawford, *et al.*, "Perspective on the current state-of-the-art of quantum computing for drug discovery applications," *Journal of Chemical Theory and Computation*, vol. 18, no. 12, pp. 7001–7023, 2022.

[6] C. Piveteau, D. Sutter, S. Bravyi, J. M. Gambetta, and K. Temme, "Error mitigation for universal gates on encoded qubits," *Phys. Rev. Lett.*, vol. 127, p. 200 505, 20 2021.

[7] K. Wan, M. Berta, and E. T. Campbell, "Randomized quantum algorithm for statistical phase estimation," *Phys. Rev. Lett.*, vol. 129, p. 030 503, 3 2022.

[8] S. Wang, S. McArdle, and M. Berta, "Qubit-efficient randomized quantum algorithms for linear algebra," 2023. arXiv: 2302.01873 [quant-ph].

[9] R. Acharya, I. Aleiner, R. Allen, *et al.*, "Suppressing quantum errors by scaling a surface code logical qubit," *Nature*, vol. 614, no. 7949, pp. 676–681, 2023.

[10] C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, "Surface code quantum computing by lattice surgery," en, *New Journal of Physics*, vol. 14, no. 12, p. 123 011, 2012, arXiv: 1111.4022.

[11] C. Gidney, "Stability Experiments: The Overlooked Dual of Memory Experiments," *Quantum*, vol. 6, p. 786, 2022, arXiv:2204.13834 [quant-ph].

[12] M. Howard and E. Campbell, "Application of a resource theory for magic states to fault-tolerant quantum computing," *Phys. Rev. Lett.*, vol. 118, p. 090 501, 9 2017.

[13] E. Campbell, "Random compiler for fast hamiltonian simulation," *Phys. Rev. Lett.*, vol. 123, p. 070 503, 7 2019.

[14] J. Eisert, D. Hangleiter, N. Walk, *et al.*, "Quantum certification and benchmarking," *Nature Reviews Physics*, vol. 2, no. 7, pp. 382–390, 2020.

[15] C. Ryan-Anderson, J. G. Bohnet, K. Lee, *et al.*, "Realization of real-time fault-tolerant quantum error correction," *Phys. Rev. X*, vol. 11, p. 041 058, 4 2021.

[16] N. J. Ross and P. Selinger, "Optimal ancilla-free clifford+t approximation of z-rotations," *Quantum Info. Comput.*, vol. 16, no. 11–12, pp. 901–953, 2016.

[17] E. H. Chen, T. J. Yoder, Y. Kim, *et al.*, "Calibrated decoders for experimental quantum error correction," *Phys. Rev. Lett.*, vol. 128, p. 110 504, 11 2022.

[18] L. Egan, D. M. Debroy, C. Noel, *et al.*, *Fault-tolerant operation of a quantum error-correction code*, 2021. arXiv: 2009.11482 [quant-ph].

[19] J. F. Marques, B. M. Varbanov, M. S. Moreira, *et al.*, "Logical-qubit operations in an error-detecting surface code," *Nature Physics*, vol. 18, no. 1, pp. 80–86, 2022.

[20] C. Ryan-Anderson, N. C. Brown, M. S. Allman, *et al.*, *Implementing fault-tolerant entangling gates on the five-qubit code and the color code*, 2022. arXiv: 2208.01863 [quant-ph].

[21] *The IBM quantum development roadmap*, en, https://www.ibm.com/quantum/roadmap, Accessed: 2023-03-28.

[22] *Our quantum computing journey*, en, https://quantumai.google/learn/map, Accessed: 2023-03-28.

[23] *Scaling IonQ's quantum computers: The roadmap*, en, https://ionq.com/posts/december-09-2020-scaling-quantum-computer-roadmap, Accessed: 2023-03-28.

[24] *The future of quantum hardware*, en, https://www.quantinuum.com/hardware/roadmap, Accessed: 2023-03-28.

Now is the time to prepare for quantum simulation of classical fluid dynamics

Johannes Mülmenstädt,[1,*] Xiangyu Li,[1] Xiaolong Yin,[2] Nathan Wiebe,[3,1,4]
Jaehun Chun,[1,5] Gregory K. Schenter,[1] and Margaret S. Cheung[1,6]

[1] *Pacific Northwest National Laboratory, Richland, WA, USA*
[2] *Petroleum Engineering, Colorado School of Mines, Golden, CO, USA*
[3] *Department of Computer Science, University of Toronto, Toronto, ON, Canada*
[4] *Canadian Institute for Advanced Research, Toronto, ON, Canada*
[5] *Levich Institute and Department of Chemical Engineering, CUNY City College of New York, New York, NY, USA*
[6] *Department of Physics, University of Washington, Seattle, WA, USA*

TOPIC

Applications, specifically numerical simulations of classical fluid dynamics and nonlinear transport problems on quantum computers

CHALLENGE

Quantum algorithms are known to handle large systems of differential equations efficiently (with computational complexity logarithmic in number of degrees of freedom) as long as they are well conditioned and linear [1–3]. For nonlinear systems such as the Navier–Stokes equations that govern fluid dynamics, the picture is less rosy; Liu *et al.* [4] have shown that efficient quantum computation for nonlinear differential equations is only possible when the ratio of nonlinear to linear terms is small. However, many important fluid dynamics applications found in fields as diverse as chemical engineering and climate change projections are characterized by strong nonlinearity.

OPPORTUNITY

Recently, we [5] and others [6] have built on the Liu et al. Carleman-linearization approach by first transforming the Navier–Stokes equations into discrete-velocity Boltzmann equations [7]. This transformation effects two favorable trade-offs. First, it replaces the turbulent nonlinearity in Navier–Stokes with a compressibility nonlinearity in the Boltzmann equation; this allows turbulent but weakly compressible flows to be efficiently linearized. Second, it trades nonlinearity for an increase in degrees of freedom, which is favorable as long as the cost of degrees of freedom is logarithmic.

What has become clear in our work is that specific nonlinear applications can be far more suitable for quantum computation than general-purpose (worst-case) complexity analysis would suggest. This is because a domain application may have specific properties (e.g., symmetries) that are not taken into account in a general-purpose complexity analysis.

This provides an opportunity to make quantum computing relevant for numerous important multiscale physics applications – e.g., climate [8–12], biological systems [13], and synthesis of advanced materials [14–18] – that are otherwise facing an intractable computational problem [19, 20]. This opportunity requires collaboration between domain scientists and quantum algorithms developers on numerous topics that were not addressed in our initial manuscript. These topics include how to treat complex boundary conditions, an open problem in quantum solvers for differential equations.

ASSESSMENT

In the near term (i.e., before quantum computers with large numbers of qubits become available), the metric of success for this work is provable upper bounds on the computational complexity (run time or qubit or gate count) of

* johannes.muelmenstaedt@pnnl.gov

quantum algorithms applied to fluid dynamics and nonlinear transport problems. In our manuscript, we have sketched a complexity analysis for homogeneous turbulent flow with periodic boundary conditions. Success for this field would be provable bounds for heterogeneous flow with complex boundary conditions.

TIMELINESS OR MATURITY

Efficient quantum computation of fluid dynamics (i.e., logarithmic scaling with resolution and domain size) has experienced breakthroughs in just the past months. Building on our recent success, a fundamentally new computational approach to fluid dynamics problems, nonlinear transport, and simulation of general multiscale phenomena now appears within reach: the efficient handling of large state spaces on quantum computers will permit simply brute-forcing the scale gap problem. Provided domain and quantum algorithms experts can come together to formulate classical physical systems in a way that plays to the strengths of quantum computers, "grand challenges" [21] in science [22] and engineering [23] will become solved problems.

---

[1] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, Phys. Rev. Lett. **103**, 150502 (2009).

[2] D. W. Berry, A. M. Childs, A. Ostrander, and G. Wang, Quantum Algorithm for Linear Differential Equations with Exponentially Improved Dependence on Precision, Communications in Mathematical Physics **356**, 1057 (2017).

[3] S. Aaronson, Read the fine print, Nature Physics **11**, 291 (2015).

[4] J.-P. Liu, H. Ø. Kolden, H. K. Krovi, N. F. Loureiro, K. Trivisa, and A. M. Childs, Efficient quantum algorithm for dissipative nonlinear differential equations, Proceedings of the National Academy of Science **118**, e2026805118 (2021).

[5] X. Li, X. Yin, N. Wiebe, J. Chun, G. K. Schenter, M. S. Cheung, and J. Mülmenstädt, Potential quantum advantage for simulation of fluid dynamics (2023), arXiv:2303.16550 [quant-ph].

[6] W. Itani, K. R. Sreenivasan, and S. Succi, Quantum algorithm for lattice boltzmann (qalb) simulation of incompressible fluids with a nonlinear collision term (2023), arXiv:2304.05915 [quant-ph].

[7] S. Chen and G. D. Doolen, Lattice boltzmann method for fluid flows, Annual Review of Fluid Mechanics **30**, 10.1146/annurev.fluid.30.1.329 (1998).

[8] K. Hasselmann, Stochastic climate models .1. theory, Tellus **28**, 473 (1976).

[9] A. J. Majda and R. Klein, Systematic multiscale models for the tropics, Journal of the Atmospheric Sciences **60**, 393 (2003).

[10] S. Bony, B. Stevens, D. M. W. Frierson, C. Jakob, M. Kageyama, R. Pincus, T. G. Shepherd, S. C. Sherwood, A. P. Siebesma, A. H. Sobel, M. Watanabe, and M. J. Webb, Clouds, circulation and climate sensitivity, Nature Geoscience **8**, 261 (2015).

[11] M. Ghil and V. Lucarini, The physics of climate variability and climate change, Reviews of Modern Physics **92**, 035002 (2020).

[12] S. Gupta, N. Mastrantonas, C. Masoller, and J. Kurths, Perspectives on the importance of complex systems in understanding our climate and climate change-the nobel prize in physics 2021, Chaos **32**, 052102 (2022).

[13] M. Bernaschi, S. Melchionna, and S. Succi, Mesoscopic simulations at the physics-chemistry-biology interface, Reviews of Modern Physics **91**, 025004 (2019).

[14] M. A. Boles, M. Engel, and D. V. Talapin, Self-assembly of colloidal nanocrystals: From intricate structures to functional materials, Chemical Reviews **116**, 11220 (2016).

[15] V. Salzmann, Bastiaan B., M. M. van der Sluijs, G. Soligno, and D. Vanmaekelbergh, Oriented attachment: From natural crystal growth to a materials engineering tool, Accounts of Chemical Research **54**, 787 (2021).

[16] J. J. De Yoreo, P. U. P. A. Gilbert, N. A. J. M. Sommerdijk, R. L. Penn, S. Whitelam, D. Joester, H. Zhang, J. D. Rimer, A. Navrotsky, J. F. Banfield, A. F. Wallace, F. M. Michel, F. C. Meldrum, H. Coelfen, and P. M. Dove, Crystallization by particle attachment in synthetic, biogenic, and geologic environments, Science **349**, aaa6760 (2015).

[17] J. J. De Yoreo, E. Nakouzi, B. Jin, J. Chun, and C. J. Mundy, Spiers memorial lecture: Assembly-based pathways of crystallization, Faraday Discussions **235**, 9 (2022).

[18] W. B. Russel, W. Russel, D. A. Saville, and W. R. Schowalter, *Colloidal dispersions* (Cambridge University Press, 1991).

[19] M. A. Sprague, S. Boldyrev, P. Fischer, R. Grout, W. I. Gustafson, Jr., and R. Moser, Turbulent flow simulation at the exascale: Opportunities and challenges workshop: August 4-5, 2015, Washington, D.C. (2017).

[20] T. Schneider, J. Teixeira, C. S. Bretherton, F. Brient, K. G. Pressel, C. Schär, and A. P. Siebesma, Climate goals and computing the future of clouds (2017).

[21] G. S. Omenn, Grand challenges and great opportunities in science, technology, and public policy, Science **314**, 1696 (2006).

[22] World Climate Research Programme, WCRP Grand Challenges (last accessed 2023–01–20), `https://www.wcrp-climate.org/grand-challenges/grand-challenges-overview`.

[23] National Academy of Engineering, *Grand Challenges for Engineering: Imperatives, Prospects, and Priorities: Summary of a Forum*, edited by S. Olson (The National Academies Press, Washington, DC, 2016).

TITLE

Optimal Quantum Circuit Compilation: Algorithms and Software

AUTHORS

Harsha Nagarajan (harsha@lanl.gov), Carleton Coffrin (cjc@lanl.gov), Los Alamos National Laboratory (LANL)

TOPIC

"Compilation"

CHALLENGE

The current resource requirements for the execution of quantum algorithms on practical applications (e.g., quantum chemistry and quanutm simulation) are daunting, requiring thousands of qubits and millions of gates [1]. Compilation of these quantum algorithms into more efficient hardware-specific Quantum Computing (QC) circuits provides a promising path to dramatically reduce the runtime (i.e., gate depth) of such computations and has received significant attention in both industrial and academic research [2]. At its core, the task of quantum compilation can be framed as how to decompose an arbitrary entangled unitary operator, into a *quantum circuit*, that is, a discrete sequence of single and multi-qubit gates that are supported by a specific quantum architecture. The first challenge in quantum compilation is that the size of unitary operator grows exponentially with the number of qubits that are consisted. Another critical challenge that characterizes quantum compilation is that specific gate operations may be forbidden, for example each qubit may not have a direct connection to all other qubits. As a consequence, certain entangling multi-qubit gates (like a CNOT gate) cannot be placed in a circuit if the target and control qubits are not physically connected [3]. Hence, implementing any arbitrary quantum circuit in such devices with hardware limitations, can lead to (a) an increase in the minimal circuit depth and (b) an increase in the total number of entangling CNOT gates, which are more time consuming than simpler non-entangling gates.

To circumvent these challenges, an important research objective is to compile compact circuits via *quantum circuit optimization*. Optimization using brute-force-type enumeration approaches with minor enhancements have been found very inefficient even on small-scale circuits [4]. More recently, there has been a significant interest to develop quick heuristic-based, first-order optimization methods to discover such compact circuits, albeit without any optimality guarantees [3, 5], and are often far from the best possible circuit (up to 50% longer circuits [6]). Some efforts also include methods for qubit routing to map a compiled circuit in to a hardware with limited connectivity of qubits [7]. However, a key challenge is that the QC community is lacking rigorous mathematical methods that can provide guarantees on the solution quality of compiled quantum circuit.

OPPORTUNITY

A key observation is that the tasks of decomposing an arbitrary unitary operator of *n* qubits, can be posed as a nonlinear discrete optimization problem [8], solution to which not only provides the quality of the compiled circuit but can also be used a feasibility verification tool based on the available gates in the elementary set and other hardware specific constraints. Initial work on this approach to quantum circuit compilation has been very successful at the scales of 2-4 qubits with circuit depths around 10; providing optimization proofs for the first time of some well known circuits and discovering new encodings for less well known gate sets [6]. However, scaling such methods to larger numbers of qubits and depth is critical to be able to observe the necessary quantum advantage over it's classical counterparts. *First*, development of tailor-made optimization (branch-and-bound-based) algorithms, incorporating physics-informed constraints (such as symmetry-breaking constraints due to gate commutation) could be critical for increased scalibility and faster convergence. *Second*, sparsity in the elementary gates could be exploited, particularly for one and two qubit gates, when placed in larger qubit circuits, which could lead to loosely-coupled smaller scale circuit compilation problems. Further these hierarchical mathematical optimization tasks could be solved using advanced decomposition algorithms and in parallel using HPC. *Third*, software development for compiler design algorithms which are capable of adapting to emerging trends in quantum programming and enable automatic bridges with complied circuits into customized hardware architectures.

ASSESSMENT

It seems possible that the proposed methods could lead to up to 50% reduction in the total depth of compiled circuits in comparison with heuristic-based methods within a few minutes on a laptop.

TIMELINESS

Recent developments of synergies between operations research and quantum computing [8, 6] will enable significant progress in rigorous development of algorithms and software for circuit compilation. The success of proposed circuit compilation methods will directly facilitate in realizing quantum advantage for algorithms in large-scale implementations. Moreover, it will provide a competitive edge for U.S. national labs in optimization methods for quantum circuit compilation when compared to similar existing capabilities at industries such as IBM, Rigetti and Google.

REFERENCES

[1] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoefler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vaschillo, "Assessing requirements to scale to practical quantum advantage," 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2211.07629

[2] A. Cho, "DOE pushes for useful quantum computing," 2018.

[3] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, "Quantum-assisted quantum compiling," *Quantum*, vol. 3, p. 140, 2019.

[4] W. N. Hung, X. Song, G. Yang, J. Yang, and M. Perkowski, "Optimal synthesis of multiple output boolean functions using a set of quantum gates by symbolic reachability analysis," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1652–1663, 2006.

[5] J.-H. Bae, P. M. Alsing, D. Ahn, and W. A. Miller, "Quantum circuit optimization using quantum karnaugh map," *Scientific reports*, vol. 10, no. 1, pp. 1–8, 2020.

[6] H. Nagarajan, O. Lockwood, and C. Coffrin, "QuantumCircuitOpt: An open-source framework for provably optimal quantum circuit design," in *SC21: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, ser. Second Workshop on Quantum Computing Software. IEEE Computer Society, 2021, pp. 55–63.

[7] G. Nannicini, L. S. Bishop, O. Günlük, and P. Jurcevic, "Optimal qubit assignment and routing via integer programming," *ACM Trans. on Quantum Computing*, vol. 4, no. 1, pp. 1–31, 2022.

[8] O. Parekh, "Synergies between operations research and quantum information science," *INFORMS Journal on Computing*, vol. 35, no. 2, pp. 266–273, 2023. [Online]. Available: https://doi.org/10.1287/ijoc.2023.1268

Photonic nanoemitters present a promising future for ideal quantum sources

Raymond Newell, [1*] Kristina Meier, [2] Duncan Ryan, [3] Jennifer Hollingsworth, [3] Charlotte Duda, [1]

[1] Materials Physics Applications - Quantum Division, Los Alamos National Laboratory
[2] Intelligence and Space Research Division, Los Alamos National Laboratory
[3] Center for Integrated Nanotechologies, Los Alamos National Laboratory
[*] corresponding author: raymond@lanl.gov

Abstract: Engineered quantum nano-emitters, such as quantum dots, carbon nanotubes, and other engineered nanostructures offer unique advantages as sources for photonic quantum information. Unlike spontaneous parametric downconversion (SPDC) sources, nanoemitters can be deterministic, true single-photon sources at almost any wavelength from visible to infrared. We discuss the importance of these attributes in the context of entanglement swapping for quantum networks.

1. Topic

Error correction and mitigation, integration

2. Challenge

Scalable quantum networks will unlock considerable new science capabilities in fundamental physics, quantum computing, cybersecurity, and beyond. Unlike classical links, which can be extended indefinitely by periodically amplifying the signal, quantum links must balance distance and channel capacity [1]. This fundamental limitation arises from the no-cloning theorem, which states that an arbitrary quantum state cannot be duplicated (i.e. amplified) with perfect fidelity [2]. Given the inevitable losses in optical fiber transmission links (≈ 0.2 dB/km) and the large loss associated with space-to-ground optical links (≈ 20 dB as in [3]), a continental-scale quantum network will require a means to coherently join multiple point-to-point links.

Entanglement swapping protocols [4–6], offer the best hope for establishing long-range quantum entanglement beyond the reach of a single link. By generating multiple sets of entangled photons, then performing Bell State Measurements (BSM) between individual photons from different sets, entanglement can be exchanged from one set to the following one, effectively doubling the distance over which entanglement is shared. This process can be repeated for even greater range until noise, attenuation or the limited efficiency of the detectors drive the success probability below acceptable limits.

3. Opportunity

The workhorse for generating entangled photon pairs has been through spontaneous parametric down-conversion (SPDC) in nonlinear crystals. SPDC photon emission is highly stochastic, delivering at-best four photon pairs per 10^6 pump photons [7]. Worse, SPDC sources have a non-zero likelihood of generating more than one pair of photons at a time, and increasing the pump power will only increase the likelihood of creating multiple pairs. Reducing the pump power to reduce the probability of multiple pairs inevitably leads to increased production of vacuum states, i.e. pulses yielding no entangled pairs.

Multiple pair production is particularly detrimental to entanglement swapping [8]. Swapping protocols can discard vacuum states via post-selection, just as is done for losses from channel attenuation. Typically, the entangling BSM does not produce a result and is simply discarded. By contrast, higher-order photon states such as multiple pairs generate a Bell State Measurement which appears valid, but is not. As a result, the error rate of a swapping operation increases dramatically. Over a chain of multiple swap operations, the rate of successful end-to-end entanglement swapping falls rapidly. Some rate can be recovered by using photon-number-resolving detectors prior to the BSM [9], but there is still no defense against producing two entangled pairs but detecting the wrong one. The only defense against such errors is a source which can generate one and only one entangled pair for each input pulse.

Quantum emitters, such as quantum dots (QDs) offer a promising path to deterministic entangled pair photon emission. Nanostructures can be engineered such that only two quanta of energy can be bound; stimulated emission (photoluminescnce) then results in exactly two photons being released. QDs are also naturally suited for deployment at scale (including space flight) due to their size and ability to generate entangled photons without bulky components that need remote stabilization.

4. Assessment

For all the promise of QDs, RD is still needed to overcome key obstacles. As a point-source emitter, a QD in an isotropic environment will be an omnidirectional emitter, sending light uniformly into 4π steradians, which couple inefficiently into distant nodes, particularly if coupling is via low numerical aperture (NA) channels such as optical fibers or free-space links. Recent advances in engineered nanostructures have addressed this limitation by resonantly coupling the QD with a beamforming structure.

5. Maturity

Researchers have developed ultra-photostable colloidal QDs capable of on-demand single-photon emission in telecommunications wavelengths at room-temperature [10]. Integration with hybrid bullseye antennas can create highly directional QD emission, particularly when performed via a novel direct-write technique using dip-pen nanolithography for deterministic placement of QDs into antenna structures. The combination of precision placement and high-quality antennas enabled demonstration of highly directional emission and record collection efficiencies of 85% into a low NA of 0.5 in free-space [11]. A next step will be to optimize fiber coupling of a QD single-photon source [12].

Once high coupling efficiency into a single-mode fiber is demonstrated, the single-photon or entangled-pair source can be integrated into virtually any network of quantum nodes connected via standard telecom fiber. These fiber-coupled, highly deterministic sources would allow for entanglement swapping at much higher rates than SPDC sources. Further scientific discovery will address the quantum emitter itself: Although the aforementioned QDs have photostable emission at telecom wavelengths at room temperature, there is still the question of whether these sources are true quantum sources, i.e., can the photons coming from one QD be indistinguishable from each other. Truly indistinguishable photons (demonstrating a Hong-Ou Mandel dip) depend on the bandwidth of the emission from a single QD, which must be sufficiently narrow. Several methods have been proposed to reduce the emission bandwidth, including composition/ geometry, operating at cryogenic temperature ($\sim$4K) and/or placing the nanoemitter into a resonant cavity structure.

References

1. M. Takeoka, S. Guha, and M. M. Wilde, "Fundamental rate-loss tradeoff for optical quantum key distribution," Nature communications **5**, 5235 (2014).
2. W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," Nature **299**, 802–803 (1982).
3. J. Yin et al., "Satellite-based entanglement distribution over 1200 kilometers," Science **356**, 1140–1144 (2017).
4. M. Zopf, R. Keil, Y. Chen, J. Yang, D. Chen, F. Ding, and O. G. Schmidt, "Entanglement swapping with semiconductor-generated photons violates bell's inequality," Physical review letters **123**, 160502 (2019).
5. R.-B. Jin, M. Takeoka, U. Takagi, R. Shimizu, and M. Sasaki, "Highly efficient entanglement swapping and teleportation at telecom wavelength," Scientific reports **5**, 9333 (2015).
6. Q.-C. Sun, Y.-F. Jiang, Y.-L. Mao, L.-X. You, W. Zhang, W.-J. Zhang, X. Jiang, T.-Y. Chen, H. Li, Y.-D. Huang *et al.*, "Entanglement swapping over 100 km optical fiber with independent entangled photon-pair sources," Optica **4**, 1214–1218 (2017).
7. M. Bock, A. Lenhard, C. Chunnilall, and C. Becher, "Highly efficient heralded single-photon source for telecom wavelengths based on a ppln waveguide," Opt. Express **24**, 23992–24001 (2016).
8. P. Dhara, S. J. Johnson, C. N. Gagatsos, P. G. Kwiat, and S. Guha, "Heralded multiplexed high-efficiency cascaded source of dual-rail entangled photon pairs using spontaneous parametric down-conversion," Phys. Rev. Appl. **17**, 034071 (2022).
9. D. Rosenberg, A. E. Lita, A. J. Miller, and S. W. Nam, "Noise-free high-efficiency photon-number-resolving detectors," Phys. Rev. A **71**, 061803 (2005).
10. S. Krishnamurthy, A. Singh, Z. Hu, A. V. Blake, Y. Kim, A. Singh, E. A. Dolgopolova, D. J. Williams, A. Piryatinski, A. V. Malko, H. Htoon, M. Sykora, and J. A. Hollingsworth, "PbS/CdS quantum dot room-temperature single-emitter spectroscopy reaches the telecom O and S bands via an engineered stability," ACS Nano **15**, 575–587 (2021). PMID: 33381968.
11. H. Abudayyeh, B. Lubotzky, A. Blake, J. Wang, S. Majumder, Z. Hu, Y. Kim, H. Htoon, R. Bose, A. V. Malko, J. A. Hollingsworth, and R. Rapaport, "Single photon sources with near unity collection efficiencies by deterministic placement of quantum dots in nanoantennas," APL Photonics **6**, 036109 (2021).
12. B. Lubotzky, A. Nazarov, H. Abudayyeh, L. Antoniuk, N. Lettner, V. Agafonov, A. V. Bennett, J. A. Hollingsworth, A. Kubanek, and R. Rapaport, "Room temperature fiber-coupled single-photon devices based on colloidal quantum dots and SiV centers in back excited nanoantennas," (2023).

Exploring strongly correlated quantum spin systems with quantum computers

Cristiano Nisoli and Coffrin Carleton
Theoretical Division, Los Alamos National Laboratory,
Los Alamos, New Mexico 87545, USA
cristiano@lanl.gov
(Dated: April 28, 2023)

Topic: Applications, models, algorithms.

Challenge: At inception, quantum annealing leveraged quantum mechanics for classical optimization tasks. However, as machine coherence increases [1–3], the realization of quantum spin systems has emerged as an even more fruitful application of this computational model [4–8]. Alternatives to the standard gate model deserve exploration, to achieve useful quantum advantage via analog quantum computing. The recent demonstration of so-called coherent quantum annealing with 1,000s of qubits [1], further expands the potential for this hardware to explore quantum system dynamics where the effects of quantum fluctuations can be carefully controlled and observed directly.

Implementing existing celebrated spin models [9–12]—or in fact new and dedicated ones—into quantum annealers will lead to observation and detection of quantum phenomena not yet observed, or not visualized directly, in experimental physics laboratories. Rather than computing or simulating quantum systems, these quantum computers allow one to simply *build* quantum systems, and experiment on them in an uniquely controlled way, with characterization down to the constitutive degree of freedom. This provides an unprecedented opportunity for understanding the physics of quantum spin systems. Theoretical physics currently abounds of interesting theoretical spin models to explore frustration, strongly correlated spins, spin liquids, fractionalized excitations, and topological matter [9–11]. However enticing, such models are generally only weak proxies for the properties of actual materials. In quantum annealers these models could be realized and experimented upon. Moreover, many more realistic models of such materials could be realized in quantum annealers. Employed in this way, quantum annealers provide an extraordinary versatile platform to explore quantum effects that are hard to find, detect, and characterize in natural materials.

Opportunity: Rather than "simulating" existing spin models with classical computers, quantum annealers can be used as Lego sets to build systems described by these models. There are many advantages to this approach.

1. Many spin model exist in a classical and quantum version. As most quantum annealers are described by the transverse quantum Ising model, they provide a unique way to control the threshold between classical and quantum behavior, which can be explored in experimental studies of coherence/decoherence.

2. Many annealing platforms allows for local control of fields acting on individual qubits. This makes a variety of constraints in experiments possible, in ways that are often impossible in real materials, as for instance the setting of boundary conditions during a quantum evolution.

3. The coupling constants among qubits can be finely modulated, which allows for the study of structural phase diagrams in these systems.

4. The recent development of so-called *reverse annealing* protocols, make it possible to initialize the qubits/spins at will to study quantum evolution from a state to another.

This opens the door to a vast variety of experiments on a rich set of models, e.g. on quantum phase transitions, spinons in a potential (by modulating the coupling constants) and their particle-statistics. Quantum annealers could be used to realize quantum spin systems in particular for the study of spinon excitations of unusual behavior. Employed this way they would provide an extraordinary versatile platform to explore quantum effects that are hard to find or to detect in natural materials.

Assessment: While *classical* spin models can certainly and usefully be realized into quantum annealers, and have been, the success of this line of work would hinge on measuring bona fide *quantum* effects, such as interference of spinon quasiparticles, or quantum entanglement. That, in turn, hinges on the coherence of new machines. Fortunately, such coherence has been increasing considerably as of late, and used to show quantum effects [1, 3, 13]. Forthcoming quantum annealing hardware will be able to exhibit more direct quantum effects, such as Friedel oscillations as interference patterns between Fermionic emergent quasiparticles. Ideally, in the future it will be possible to realize into quantum annealers spinon quasiparticles, like hard-core bosons, fermions and perhaps anyons, which could be individually controlled and characterized. This turn would prove revolutionary and bring to fruition in artificial quantum realization a couple of decades of theoretical work on quantum spin liquids and other strongly correlated spin systems.

Timeliness or maturity: This position is timely because the idea of deliberately designing interacting *classical* binary spin systems and building them into a variety of platforms is quite mature [12, 14–17]. Such platforms have ranged from interacting magnetic nano-islands [15, 16], to trapped colloids or vortices in nano-patterned superconductors [17–19], or skyrmions in magnets [20] or liquid crystals [21], macroscopic magnets [22, 23], and even to mechanical metamaterials [24, 25]. They have been used to generate materials of exotic and often pseudoparadoxical classical behaviors, such as entropy-driven order [26]. This has required an evolution in creativity, adaptation to specific platforms, and ductility in design that can be transported to a quantum annealing platform. Indeed, while the design and realization of artificially interacting *classical* spin systems of desired unusual phenomenology is a well established effort, quantum annealers can open a path to their *quantum* phenomenology, where opportunity of new insights are considerably higher.

Recent works have demonstrated implementation of classical models into quantum annealers [6–8]. Even at a classical level, they allowed for experiments that would be impossible on other platforms, *e.g.* individual characterization of magnetic monopoles and their entropic screening [6] or control between topologically-trivial and non-trivial dynamics [7]. Though quasiclassical, already these works have shown quantum effects, by activating spinons kinetics via quantum fluctuations [6, 7]. Recent work has also demonstrated enough coherence to reveal signatures of the quantum Kibble-Zurech mechanism and critical quantum dynamics [1, 2].

The confluence of an aggregated experience in developing spin models of unusual properties at the classical and quantum level, of the successful effort in realizing a variety of unusual models in various classical platforms, of early works translating these models into annealers, and of recent results in coherence, opens now a path toward the realization of *quantum* spin models of exotic properties into quantum annealers.

---

[1] A. D. King, S. Suzuki, J. Raymond, A. Zucca, T. Lanting, F. Altomare, A. J. Berkley, S. Ejtemaee, E. Hoskinson, S. Huang, *et al.*, Nature Physics **18**, 1324 (2022).
[2] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud, *et al.*, Nature , 1 (2023).
[3] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Science **374**, 1242 (2021).
[4] A. D. King, J. Carrasquilla, J. Raymond, I. Ozfidan, E. Andriyash, A. J. Berkley, M. Reis, T. Lanting, R. Harris, F. Altomare, K. Boothby, P. I. Bunyk, C. Enderud, A. Fréchette, E. M. Hoskinson, N. Ladizinsky, T. Oh, G. Poulin-Lamarre, C. Rich, Y. Sato, A. Y. Smirnov, L. J. Swenson, M. H. Volkmann, J. Whittaker, J. Yao, E. Ladizinsky, W. Mark, J. P. Hilton, and M. H. Amin, Nature **560**, 456 (2018).
[5] A. D. King, J. Raymond, T. Lanting, S. V. Isakov, M. Mohseni, G. Poulin-Lamarre, S. Ejtemaee, W. Bernoudy, I. Ozfidan, A. Y. Smirnov, *et al.*, Nature communications **12**, 1 (2021).
[6] A. D. King, C. Nisoli, E. D. Dahl, G. Poulin-Lamarre, and A. Lopez-Bezanilla, Science **373**, 576 (2021).
[7] A. Lopez-Bezanilla, J. Raymond, K. Boothby, J. Carrasquilla, C. Nisoli, and A. D. King, Nature communications **14**, 1105 (2023).
[8] A. Lopez-Bezanilla and C. Nisoli, Science Advances **9**, eadf6631 (2023).
[9] A. Kitaev, Annals of Physics **321**, 2 (2006).
[10] L. Balents, Nature **464**, 199 (2010).
[11] L. Savary and L. Balents, Reports on Progress in Physics **80**, 016502 (2016).
[12] M. J. Morrison, T. R. Nelson, and C. Nisoli, New Journal of Physics **15**, 045009 (2013).
[13] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud, E. Hoskinson, and et. al., "Quantum critical dynamics in a 5000-qubit programmable spin glass," (2022), arXiv:2207.13800.
[14] C. Nisoli, R. Moessner, and P. Schiffer, Reviews of Modern Physics **85**, 1473 (2013).
[15] S. H. Skjærvø, C. H. Marrows, R. L. Stamps, and L. J. Heyderman, Nature Reviews Physics , 1 (2019).
[16] P. Schiffer and C. Nisoli, Applied Physics Letters **118**, 110501 (2021).
[17] A. Ortiz-Ambriz, C. Nisoli, C. Reichhardt, C. J. Reichhardt, and P. Tierno, Reviews of Modern Physics **91**, 041003 (2019).
[18] Y. Han, Y. Shokef, A. M. Alsayed, P. Yunker, T. C. Lubensky, and A. G. Yodh, Nature **456**, 898 (2008).
[19] A. Libál, C. Reichhardt, and C. J. Olson Reichhardt, Phys. Rev. Lett. **97**, 228302 (2006).
[20] F. Ma, C. Reichhardt, W. Gan, C. O. Reichhardt, and W. S. Lew, Physical Review B **94**, 144405 (2016).
[21] A. Duzgun and C. Nisoli, Physical Review Letters **126**, 047801 (2021).
[22] P. Mellado, A. Concha, and L. Mahadevan, Physical review letters **109**, 257203 (2012).
[23] H. A. Teixeira, M. F. Bernardo, M. D. Saccone, F. Caravelli, C. Nisoli, and C. I. de Araujo, arXiv preprint arXiv:2112.06058 (2021).
[24] A. S. Meeussen, E. C. Oğuz, Y. Shokef, and M. v. Hecke, Nature Physics **16**, 307 (2020).
[25] C. Merrigan, C. Nisoli, and Y. Shokef, Physical Review Research **3**, 023174 (2021).
[26] H. Saglam, A. Duzgun, A. Kargioti, N. Harle, X. Zhang, N. S. Bingham, Y. Lao, I. Gilbert, J. Sklenar, J. D. Watts, *et al.*, Nature Physics , 1 (2022).

Full-Stack Codesign: Crossing Abstraction Layer Boundaries for Algorithm- and Application-Informed Optimizations to Performance of NISQ-era Quantum Computers

Ravi Naik*, Kasra Nowrouzi*, ☨

QuIST Group, AMCR Division, Lawrence Berkeley National Laboratory

Topics: codesign and integration, compilation, error mitigation, models, applications, algorithms

Challenge:

Several abstraction models have been proposed in recent years to divide up the challenges of building a full-stack quantum computer [e.g. 1], due to a lack of easy access to quantum computing platforms, among other factors. Technical understanding and language can vary for research at different layers of abstraction, leading to difficulties in cross-layer optimization. While the artificial layers of abstraction allow different segments of the QIS community to contribute to the field, they add complexity and rigidity. This results in reduced performance and efficiency in execution of quantum and hybrid algorithms. Furthermore, partly due to incentive structures in the industrial/commercial landscape, most entities who develop and deploy full stack quantum computers prioritize error correction and system scaling, aimed at large-scale universal fault-tolerant computing, instead of optimizing performance and efficiency of NISQ-era quantum computers. This is detrimental to the advancement of NISQ-era explorations of quantum and hybrid algorithms and applications, and associated enabling technologies. These explorations will be vital to the development of pathways to what will become large-scale industrial fault-tolerant computing.

Opportunity:

Quantum computers at DOE National Labs are ideally positioned to address fundamental science and technology challenges on integrated full-stack systems. At DOE National Labs and Testbeds, such as the Advanced Quantum Testbed at LBNL, full transparent access to all layers of the stack has allowed for integration of knowledge from a variety of technical expertise. This collaborative research approach has resulted in groundbreaking work on characterization and mitigation of noise and errors, compilation tools to bring more demanding experiments within reach of existing NISQ hardware platforms, and the application of all such tools to enable state-of-the-art simulations of nature on quantum computers using both qubit- and qutrit-based quantum processors.

Some of these projects have involved breaking the aforementioned abstraction layer boundaries to increase performance [2, 3], demonstrating initial explorations of this approach. To more directly enable this line of research, an LDRD project at Berkeley Lab led by Kasra Nowrouzi, a co-author of this position paper, has engaged in codesign of quantum and hybrid algorithms with control systems. We are now at a point, given the experiences described above, where codesign of the entire full stack of quantum computing is within reach. This would include traditionally distinct domains such as algorithms, tools for compilation, optimization, and transpilation, control systems (including hardware, firmware, gateware, software), and quantum processor architecture. Examples include:

1. Codesigning control systems with algorithms to improve execution efficiency, and to enable the execution of more advanced algorithms,

* Equal Contribution

☨ Corresponding Author: knowrouzi@lbl.gov

2. Codesigning compilation and optimization tools with control systems to implement and take advantage of novel control capabilities,
3. Codesigning noise and error mitigation tools to learn from specifically-designed quantum processors to study and mitigate noise,
4. Codesigning quantum processors with algorithms for application-specific computing,
5. Investigate sources of inefficiency and inaccuracy using application benchmarking and time profiling to identify further opportunities for codesign

This challenge is an ideal fit to the capabilities and organization of National Labs, due to the collocation of multidisciplinary teams of researchers involved in research on various parts of the stack, and access to related resources and expertise, such as classical supercomputing Facilities. Furthermore, DOE Testbeds, having established themselves at the nexus of collaborative research between Academia, Industry, and National Labs, are uniquely positioned to implement algorithms- and application-informed codesign of the full-stack.

Assessment:

The aim of this full-stack codesign approach will be to implement near-term algorithm- and application-informed optimizations to the performance of NISQ-era quantum computers. Assessment would include execution of algorithms and applications of interest to the DOE QIS community, through engagement with researchers involved in materials science, quantum chemistry, physics, and data analysis and optimization. Benchmarking methods, including application benchmarking, would assess the performance of individual layers and the full stack. In addition to improving fidelities and accuracy for quantum and hybrid experiments, efficient execution of algorithms is an equally important orthogonal metric, requiring improvements in order to accelerate exploration of the space of near-term algorithms of interest to the community. In the longer term, this increased efficiency can also enable faster execution of later-stage algorithms for practical quantum advantage before fault-tolerance. Thus, experiments will be profiled to assess time spent in various layers for the purpose of implementing cross-layer optimizations to improve overall execution efficiency.

Timeliness or Maturity:

Significant work has been done over the last few years to get us to where we are today, where we have gained a more detailed understanding of the challenges and opportunities associated with distinct layers of the quantum computing stack, and have made improvements to these individual layers. At the same time, error-corrected, fault-tolerant quantum computing is still many years away; and progress in development of near-term algorithms and applications will in part depend on an optimal full-stack. The present moment is thus the ideal time to engage in application-informed full-stack codesign of quantum computing platforms, to increase the performance of existing NISQ platforms and accelerate the transition to fault-tolerance.

References:

[1] X. Fu, L. Riesebos, L. Lao, C. G. Almudever, F. Sebastiano, R. Versluis, E. Charbon, and K. Bertels. 2016. A heterogeneous quantum computer architecture. In Proceedings of the ACM International Conference on Computing Frontiers (CF '16). Association for Computing Machinery, New York, NY, USA, 323–330.

[2] E. Younis, K. Sen, K. Yelick and C. Iancu, "QFAST: Conflating Search and Numerical Optimization for Scalable Quantum Circuit Synthesis," 2021 IEEE International Conference on Quantum Computing and Engineering (QCE), Broomfield, CO, USA, 2021, pp. 232-243

[3] Hashim, A., Rines, R., Omole, V., Naik, R. K., Kreikebaum, J. M., Santiago, D. I., ... & Gokhale, P. (2022). Optimized SWAP networks with equivalent circuit averaging for QAOA. *Physical Review Research*

Quantum abstract machines without circuits: the need for higher algorithmic expressiveness

Santiago Núñez-Corrales[1,2] – nunezco2@illinois.edu
[1]National Center for Supercomputing Applications
[2]Illinois Quantum Information Science and Technology Center
University of Illinois Urbana-Champaign

Topics: foundations of quantum computation; design and analysis of established and novel abstract quantum computing and programming models

Challenge: Defining what computation means requires understanding what resources are available and how these can be marshalled to accomplish tasks we identify as computing at an abstract level. Having acquired this knowledge in advance partially explains the successful evolution of classical computing technologies, in which universal Turing machines and lambda calculus delineated hardware requirements sufficiently well. Later these were satisfied through the match between bi-stable digital electronics and Boolean algebra which resulted in finite models of arithmetic in use today. To reiterate: the core theoretical models of computation via abstract machines were readily available well in advance of the transistor, which made algorithmic development agnostic to hardware details even at VLSI scales. We characterize here the information flow going from theory to hardware as *top-down*, and originated from questions about provability of theorems in (integer) arithmetic.

Quantum computing inverts the paradigm described above, going *bottom up*. All formulations of quantum Turing machines, quantum random access machines (*1*, *2*), (*3–5*) and quantum λ-calculus (*6*, *7*) so far prescribe the execution of gates on primitive quantum steps –primitive in the sense of mapping directly to quantum hardware- thus conflating the abstract concept of performing computation with the concrete task of building circuits. Three reasons seem to largely explain the latter: (a) the original intent involved modeling physical processes as computation at the most fundamental level (and viceversa), (b) reasoning about small quantum systems is more feasible that for large ones, and (c) building experiments that benefit from decades of research in quantum science was directly attainable.

Despite the ongoing quantum computing revolution and substantial funding dedicated to it, progress the number of quantum algorithms is scant. Going from gates and qubits to a general way to create new ones remains elusive. In a hypothetical situation in which classical computing followed a similar path, defining and understanding computation from the basis of bi-stable elements and Boolean algebra would have been substantially harder; the fact that quantum devices offer a much larger array of resources for computation (*8*) compounds our difficulties even more. To put it bluntly, quantum information –and by extension quantum mechanics (*9*)- does not provide a sufficient basis to understand quantum computation in abstract terms.

Opportunity: To reap the expected benefits of quantum computing, abstract machine models capable of facilitating the development of algorithms at higher conceptual levels are direly needed. Regardless of how these are specified –e.g., instructions vs functions- it is clear that outcome of such exercise will produce composable procedural abstractions: entities that operate well beyond Hilbert spaces and their transformations, that can be combined to produce useful generative effects, and that provide a denotational semantics of future quantum programming languages which makes no explicit reference to circuits. In particular, finding an analogue to finite models of integer arithmetic for quantum computing would allow specifying new classes of quantum abstract machines in terms of their (possibly discrete) transformations without worrying about hardware details of any sort, and then delegate hardware details to computer architecture designers and, later, to specialized hardware compilers.

Succinctly: composable procedural abstractions are essential for the sustained development and discovery of new quantum algorithms; these cannot be found amidst quantum circuits. Creating the funding and research context around this challenge will produce a “quantum jump” in the way we characterize computational problems (*10*), and consequentially in the number of applications for which quantum computing yields a concrete advantage.

Assessment: At present, some of the ingredients needed to develop high-level quantum abstract machines of the sort we look for seem to be present. These include: the relation between Clifford algebras and quantum field theory (*11*), the ability to define abstract machines using geometrical algebra (*12*), symmetries present in quantum random walks with implications for quantum automata (*13*), and higher algebra formulations of quantum field theory (*14–16*). However, they are insufficient for the task in their present form. It is not immediately obvious how to arrange them to produce new theories and abstract machines directly at a high level, which other ones are missing, or whether the resulting theories will hold once the limits of entanglement and other quantum resources are more deeply explored.

However, preliminary research and practice in quantum programming suggests a possible route: privileging those theories where combinators arise, and where certain patterns in circuit-building can be abstracted away as *quantum motifs* where details such as the number of qubits, specific gates and how these can be optimized disappear. Combinators proved fundamental in early days of classical computing to capture meta-patterns that simplify reasoning about how we construct algorithms and data structures. This work requires investing in interdisciplinary teams that include theoretical physicists, mathematicians, logicians, and computer scientists. While the investment for a single team looks modest in comparison to experimental work in quantum computing, the difficulty of the task calls for sufficient funding across relevant agencies to maximize the overall surface of attack through multiple research teams, thus maximizing the probability of finding good candidate theories and abstract machines.

Timeliness: Predicted timelines in quantum hardware manufacturers suggest an upcoming era of *very large quantum scale integration* (VLQSI). Devising and implementing large algorithms will become rapidly infeasible for humans. Given the level of investment in quantum computing, the capabilities of devices for theory testing, the increasing pressure to move from proof of principle to applications, and how limited progress on this issue has been achieved, the urgency of this challenge cannot be understated.

References

1. E. Knill, *arXiv preprint arXiv:2211.02559* (1996).
2. Q. Wang, M. Ying, *Journal of Computer and System Sciences* **131**, 13–63 (2023).
3. D. Deutsch, *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **400**, 97–117 (1985).
4. R. P. Feynman, *Found. Phys.* **16**, 507–532 (1986).
5. S. Guerrini, S. Martini, A. Masini, *Applied Sciences* **10**, 5551 (2020).
6. P. Selinger, B. Valiron, *et al.*, *Semantic techniques in quantum computation*, 135–172 (2009).
7. N. Botö, F. Forslund, *arXiv preprint arXiv:2303.17399* (2023).
8. E. Chitambar, G. Gour, *Reviews of modern physics* **91**, 025001 (2019).
9. J. Bub, *Foundations of Physics* **35**, 541–560 (2005).
10. S. Chapman, G. Policastro, *The European Physical Journal C* **82**, 128 (2022).
11. J. Baugh, D. R. Finkelstein, A. Galiautdinov, H. Saller, *Journal of Mathematical Physics* **42**, 1489–1500 (2001).
12. R. Schott, G. Stacey Staples, *Advances in applied Clifford algebras* **20**, 121–140 (2010).
13. P. Arnault, *Physical Review A* **106**, 012201 (2022).
14. H. Halvorson, M. Müger, *arXiv preprint math-ph/0602036* (2006).
15. G. Basti, A. Capolupo, G. Vitiello, *Progress in Biophysics and Molecular Biology* **130**, 39–52 (2017).
16. M. Benini, A. Schenkel, L. Woike, *Letters in Mathematical Physics* **109**, 1487–1532 (2019).

Causality Discovery in Biomedicine as a Killer Application for Quantum Computing

Nicholas Chia[1*], Nikita Agarwal[1] – {Agarwal.Nikita}@mayo.edu
Santiago Núñez-Corrales[2,3] – nunezco2@illinois.edu
[1]Mayo Clinic; [2]National Center for Supercomputing Applications, UIUC
[3]Illinois Quantum Information Science and Technology Center, UIUC
*Corresponding author

Topic: applications, algorithms

Challenge: A systemic issue in the machine learning space is the use of correlational models to aid decision-making which results in 87% of models never making it beyond the research phase[1]. Causal AI is considered a key step in creating machines which are able to reason and make decisions like humans do (*1*), which does not seem feasible with other AI methods, yet shows significant challenges to operate computationally at scale. Developing technologies to detect and characterize causality at scale can thus help us create a society where machines make decisions that humans are able to understand and scrutinize.

Some key lessons on the importance of causality come from stories enshrined in medical history. One of the most infamous examples comes from Nobel Prize winner Sidney Farber (*2*). Farber noticed a correlation between leukemia and several nutritional deficiencies related to vitamin B deficiency. He then reasoned that the vitamin B deficiency must be causing progression of leukemia. Based on this logic tried treating childhood leukemias with B-vitamins, resulting in his patients dying at greatly accelerated rates. Despite this set back, Farber quickly realized he had gotten the direction of causality reversed. It was the presence of vitamin B that was causing the progression of leukemia and not it's absence. Farber then went on to try B-vitamin antagonists, which then resulted in one of the first major success stories in the fight against cancer.

In Farber's case, correlation quickly lead to an understanding of causation. But what if we have a more complex causal hypothesis with more subtle effects? For example, what if we want to understand the causal role of diet and environmental exposures in cancer? We don't have to imagine this as many real-world efforts dissect this problem. Before the recent announcement from the WHO regarding the recognition of red meat as a carcinogen, a working group "considered more than 800 different studies on cancer in humans" (*3*).

Both examples point to a stark dichotomy found all over biomedical research: the world is structured but contingently and largely complex. Causality-assisted machine learning allows us to uncover the underlying structure –i.e., a cause-and-effect relationship between features of a problem- but not without penalty. First, increasing the number of features (i.e., nodes in a causal network) under consideration drastically multiples the need for compute power. For instance, estimating counterfactual bounds in an algorithmic recourse (*4*) in fully (or partially confounded) settings is exponentially expensive. Second, applying causal reasoning requires certain assumptions about the underlying relationships in our data, and many of those assumptions will remain untestable; for those in which these are ethically testable, combinatorics may make it infeasible in practice. Third, unknown confounding refers to the inability to ascertain whether other factors explain effects present in (observational) data; testing unconfoundedness can also lead to a combinatorial explosion. Clearly, new methods are needed here.

Opportunity: The properties of causal discovery and inference in biomedicine appear to match well the expected advantages brought forth by quantum computation and potentially enable advances with significant clinical impact. We believe the ingredients exist to explore algorithms to (a) encounter when causality and

[1]See: https://d2iq.com/blog/why-87-of-ai-ml-projects-never-make-it-into-production-and-how-to-fix-it.

causal effects are not identifiable and (b) perform causal inference when it is identifiable. Broadly speaking, solving tasks involving causal inference boils down to either an optimization problem, a combinatorics problem, or a combination of both. For instance, if it is an optimization problem, algorithms such as QAOA (*5*) and other algorithms have shown promise of quantum advantage. The problem of unidentifiability, whether functional or induced by confounders, is a global condition of the composition of a model and a dataset, characterizable as a global function. Following the reasoning of algorithms such as Deutsch-Jozsa, it is likely to be successful in finding the assumptions on quantum oracles or methods with quantum oracles to make unidentifiable models identifiable.

Solving problem (a) may lead to a unique solution to the problem at hand rather than finding an equivalence set of solutions and solving problem (b) should allow estimating the probability of a successful characterization as a function of the number of times the algorithm would need to be run, possibly with at least polynomial speedup, and ideally with exponential gains. Given the classical state of affairs in causal discovery and inference, even polynomial speedups can increase the feasibility of hard problems in biomedicine and clinical practice.

Assessment: The relationship between quantum computation and causality is well-known (*6*), as well as how quantum causality itself is richer than its classical counterpart (*7*). Experiments on causal relations between optical modes indicate it is possible to extract causal structure from observations only (*8*). Two quantum machine learning algorithms have been reported as well for causal discovery in knowledge graphs (*9*), using both variational (quadratic speedup) and quantum tensor-based (exponential speedup) algorithms. Causal discovery for physical processes has also been explored through quantum causal unraveling (*10*), a greedy algorithm for quantum entropic causal discovery (*11*), and the user of Grover search for causal generative models in genomics (*12*). Finally, applying causal discovery algorithms to entangled quantum systems has revealed significant subtleties arising from quantum causality (*13*), which may also impact classical applications.

Demonstrating proof of concept in this area is within reach due to the increasing availability of Noisy Intermediate-Scale Quantum (NISQ) hardware platforms with a larger number of qubits. We believe doing so requires (a) performing theoretical work to either adapt existing encodings and algorithms to capture causality in a way that maximally exploits problem structure or develop new ones, (b) constructing synthetic reference datasets representative of biomedical problems suitable for NISQ experimentation, (c) extending the work to increasingly complex problems guided by biomedical research needs. All these are feasible problems given sufficient dedicated funding across relevant agencies.

Timeliness or maturity: We believe that all elements to perform this kind of research are currently present. We estimate: 2y – proof of concept with exact solutions for at most 5 causal factors; 4y – application to clinical trials with at most 10 factors; >10y – application to constrained multifactorial models that include exposomics, a key emerging field.

References

1. R. Heil *et al.*, *TATuP-Zeitschrift für Technikfolgenabschätzung in Theorie und Praxis/Journal for Technology Assessment in Theory and Practice* **30**, 30–36 (2021).
2. S. Mukherjee, *The emperor of all maladies: a biography of cancer* (Simon and Schuster, 2011).
3. V. Bouvard *et al.*, *The Lancet Oncology* **16**, 1599–1600 (2015).
4. S. Venkatasubramanian, M. Alfano, presented at the Proceedings of the 2020 conference on fairness, accountability, and transparency, pp. 284–293.
5. E. Farhi, J. Goldstone, S. Gutmann, *arXiv preprint arXiv:1411.4028* (2014).
6. Č. Brukner, *Nature Physics* **10**, 259–263 (2014).
7. J.-P. W. MacLean, K. Ried, R. W. Spekkens, K. J. Resch, *Nature communications* **8**, 15149 (2017).
8. K. Ried *et al.*, *Nature Physics* **11**, 414–420 (2015).
9. Y. Ma, PhD thesis, lmu, 2020.
10. G. Bai, Y.-D. Wu, Y. Zhu, M. Hayashi, G. Chiribella, *npj Quantum Information* **8**, 69 (2022).
11. M. A. Javidian, V. Aggarwal, F. Bao, Z. Jacob, presented at the Quantum Information and Measurement, F2C–3.
12. A. Sarkar, Z. Al-Ars, K. Bertels, *Applied Sciences* **11**, 2696 (2021).
13. C. J. Wood, R. W. Spekkens, *New Journal of Physics* **17**, 033002 (2015).

Quantum algorithms for dynamical systems: ASCR position paper

Matteo Lostaglio, David Jennings (to present)*, Sam Pallister (PsiQuantum)
Robert Lowrie, Yigit Subasi, Andrew Sornborger (Los Alamos National Laboratory)

I. Topic - Applications, Algorithms, Co-Design

II. Introduction

Dynamical systems, such as those in plasma physics and fluid-dynamics, are major challenges for classical simulation. A substantial fraction of today's high-performance supercomputing is devoted to the resolution of these problems. Fluid-dynamics and plasma physics projects, for example, make up around 40% of the total core hours allocated within the INCITE 2023 program running on the high-performance supercomputers administered by the Department of Energy [1]. In other words, these problems are not only extremely relevant, but they also face serious computational bottlenecks on today's largest supercomputers. It is therefore a pressing question to determine the potential of quantum computing in this area, and to expand both the scope and importance of quantum computers under development.

While quantum computation has been extensively studied for its significant applications in quantum chemistry and quantum physics, far less is understood about its impact on the study of high-dimensional classical dynamical systems. Very recently there has been a range of significant advancements in this direction [2–5]. At least two quantum algorithms have shown evidence that useful speedups can be achieved for fluid simulations and related classical systems of equations [2, 4]. Depending on the algorithmic details of the problem solved, a quadratic speedup is attainable and there is evidence that an exponential speedup may also exist. These results indicate that quantum algorithms for classical systems are an important possible DOE mission-relevant area of application for future, fault-tolerant quantum computers.

III. Challenge

The simulation of quantum systems is expected to be a highly disruptive application of quantum computing. However, currently the simulation of classical dynamical systems on quantum computers, in particular non-linear ones, is significantly less well-understood.

The analogy with quantum chemistry is enlightening in this respect. In the last few years considerable effort has been devoted to identifying suitable target problems and parameter regimes; to embed into the quantum algorithm cost reductions due to physical insights and the use of cutting-edge classical techniques; and to gauge progress in relation to classical benchmarks and algorithms. Progress in each of these directions is pivotal to the process of making quantum chemistry algorithms feasible on early to mid-term fault-tolerant quantum computers. Quantum algorithms groups around the world have been targeting gold-standard problems (in particular, simulation of FeMoco [6]) with the help of domain experts and progressively found more efficient realizations of core subroutines of the algorithm. In the last six years, resource estimates have come down by six orders of magnitude [7].

In order to provide such specificity, it is not enough to focus on asymptotic scaling; even algorithms with "good scaling" may require a quantum computer of impractical physical size and power requirements. In the context of quantum chemistry, algorithms have instead been compiled in detail, to the point that we can now give exact resource estimates of the running costs on fault-tolerant quantum computers given certain architecture choices. This also means having tools to match algorithmic and compilation advances with hardware projections, obtaining valuable estimates of when a certain problem will be runnable on a quantum computer of a given specification. This compilation pipeline does not exist for quantum algorithms for simulating classical dynamical systems. We therefore do not have rigorous, quantitative estimates of when such algorithms will be valuable to the DOE.

We believe that the construction of an ecosystem similar to the one just described will be crucial to progress along a trajectory where we try to make, for example, nonlinear plasma and fluid-dynamics problems a serious target for mid-term quantum computers. While nonlinear problems have recently seen a flurry of new results, algorithms are still few, detailed running costs are unavailable and there is no obvious roadmap towards tackling increasingly challenging problems. It is to discuss the opportunities lying behind each of these three challenges that we now turn.

IV. Opportunity

We identify three core components to tackling the above challenge: (i) the development of algorithmic reduction pipelines; (ii) the development of tools for end-to-end compilation and resource estimation; (iii) the input from experts in multiple domains to identify "gold-standard" problems and to contribute to importing relevant classical techniques. Let us briefly discuss each.

Algorithmic reduction pipelines involve bridging the gap between the mathematical formulation of nonlinear differential equations and problems natively solved on quantum computers. There are several routes to do so. The common feature among them is that the nonlinear problem is embedded into a linear one – a move that is not advisable classically due to the curse of dimensionality, but which has much better prospects on a quantum computer where we can manipulate large vectors encoded as quantum states. The linear problem, suitably truncated and discretized, is eventually reduced

* djennings@psiquantum.com

to Hamiltonian simulation, often [2, 8] via quantum linear solver algorithms. Convergence guarantees of the reduction methods and bounding the relevant algorithmic parameters of the 'reduced' problems (e.g., the condition number of the associated linear system) as a function of the parameters of the original nonlinear equation are key technical challenges in this analysis.

Secondly, compilation and resource estimation, even in more established domains such as quantum chemistry, continues to be an "artisanal" process, relying on algorithm constructions and resource bookkeeping done by hand. This need not be the case. Nascent software efforts [9, 10] have begun to automate some of this process, including both algorithm construction and resource estimation. Establishing a suite of general-purpose algorithms and subroutines tailored to quantum simulation of dynamical systems would enable the field to rapidly iterate through new systems, instances and algorithm designs.

Finally, making progress towards viable and useful quantum applications in this space will require the input of experts from multiple domains: Hamiltonian simulation theory, quantum linear solver methods, classical dynamical systems, and the theory of differential equations. We must provide conditions for researchers to bring their expertise to bear in a new setting, and tackle significant problems in an emerging field. Should we be successful in drawing together a diverse community of researchers around this domain, we will engender an ecosystem that facilitates new collaborations, more easily distils emerging techniques, and rapidly identifies targets that are: (a) classically intractable; (b) amenable to near/medium term quantum simulation; and (c) of real-world significance.

V. Assessment

Our three core components outlined in the previous section can be assessed as follows:

1) Establishment of a set of "gold-standard" systems and instances. These instances must be (i) strategically relevant for the DOE and the community at large; (ii) challenging to simulate on large-scale conventional hardware, and (iii) have quantifiable resource estimates that we believe are in reach of near- or medium-term quantum hardware.
2) Implementation of suites of general-purpose algorithms and subroutines instantiated in software, connected to a robust resource estimation pipeline that is capable of producing both logical- and physical-level resource estimates.
3) Development and support of the creation of an interdisciplinary research community around this topic, in particular one that is capable of providing input to points (i) and (ii) above, and that can contribute to novel algorithmic advances in cases where optimizations from the classical domain can be ported to the quantum one.

VI. Timeliness and maturity

Until recently, three prerequisite components were lacking: (i) established results on asymptotic speedup for simulation of dynamical systems; (ii) case studies with which to emulate the compilation pipeline; and (iii) software tools to automate algorithm composition and resource estimation. While all are nascent, there is now a firm footing in each.

Additionally, quantum hardware has improved significantly in the past three years. Multiple early implementations of error correcting codes have been implemented demonstrating logical qubits with extended lifetimes. Roadmaps from vendors point to systems that are capable of implementing meaningful fault-tolerant quantum computations by the end of this decade, and the architectures in which those machines operate are evolving rapidly. It is imperative that we develop applications that are capable of automatically tracking hardware and architectural advances and are ready to fully exploit the technology when it is delivered.

Lastly, the strategic relevance of this family of computational problems cannot be overstated. Defining a viable path to quantum advantage now has the potential of broad impact across myriad scientific disciplines, spanning from fundamental science to industry, and can help drive strategy and investment decisions across the ecosystem.

LANL unclassified release designation: LA-UR-23-24332.

References

[1] Department of Energy, *INCITE program 2023*, (2023) doeleadershipcomputing.org.

[2] J.-P. Liu, H. Ø. Kolden, H. K. Krovi, N. F. Loureiro, K. Trivisa, and A. M. Childs, "Efficient quantum algorithm for dissipative nonlinear differential equations", Proceedings of the National Academy of Sciences **118** (2021).

[3] S. Jin, N. Liu, and Y. Yu, "Quantum simulation of partial differential equations via schrodingerisation: technical details", arXiv preprint arXiv:2212.14703 (2022).

[4] X. Li, X. Yin, N. Wiebe, J. Chun, G. K. Schenter, M. S. Cheung, and J. Mülmenstädt, *Potential quantum advantage for simulation of fluid dynamics*, 2023.

[5] R. Babbush, D. W. Berry, R. Kothari, R. D. Somma, and N. Wiebe, "Exponential quantum speedup in simulating coupled classical oscillators", arXiv preprint arXiv:2303.13012 (2023).

[6] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, "Elucidating reaction mechanisms on quantum computers", Proceedings of the National Academy of Sciences **114**, 7555–7560 (2017).

[7] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, "Even more efficient quantum computations of chemistry through tensor hypercontraction", PRX Quantum **2** (2021).

[8] H. Krovi, "Improved quantum algorithms for linear and nonlinear differential equations", Quantum **7**, 913 (2023).

[9] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoefler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vaschillo, *Assessing requirements to scale to practical quantum advantage*, 2022.

[10] DARPA, *Quantum benchmarking*, (2023) darpa.mil/program/quantum-benchmarking.

System Software Stack for Diverse Quantum Computing Technologies

Tirthak Patel
Rice University, Houston, TX
tirthak.patel@rice.edu

Devesh Tiwari
Northeastern University, Boston, MA
d.tiwari@northeastern.edu

I. Topic: Compilation

Due to its potential to revolutionize various industries by solving classically complex problems, quantum computing has undergone rapid development. However, much of the focus has been on one type of quantum computing technology – superconducting qubits – especially in the context of the quantum compilation and software stack [3]. In this paper, we discuss the challenges and opportunities of developing quantum software compilation and stack solutions and extending exisiting solutions for other technologies like photonic qubits, trapped ion qubits, and neutral atom qubits.

II. Challenge

Most of the focus in quantum computing has been on superconducting qubits, which have unique properties such as high scalability, limited qubit connectivity and low coherence times [4]. The current quantum software stack is designed to minimize the effect of hardware noise on current quantum computers by optimizing for these properties, and it may not be suitable for other quantum computing technologies. While superconducting qubits are the most advanced quantum computing technology so far – due to the ease of porting classical semiconductor technology to quantum computing –, other under-development technologies are also promising [7].

Photonic qubits, trapped ion qubits, and neutral atom qubits have different properties that require tailored solutions for software and compiler stack design. For example, photonic qubits have greater output usage flexibility and long coherence times, but they also have low a relatively lower operation fidelity; this requires error mitigation strategies that are different than the ones deployed for superconducting qubits. *Developing software solutions that optimize for these properties is challenging, especially while ensuring that the user is agnostic of the technology-specific optimizations for greater accessibility.*

III. Opportunity

To address the challenges of developing software solutions for other quantum computing technologies, there is a need for new tools and techniques that can optimize for their specific properties. Developing new software-based error mitigation techniques for technologies like trapped ion qubits and neutral atom qubits can help improve their fidelity and make them more suitable for practical applications. Moreover, developing a modular and portable software stack would enable scientists and users to easily switch between different quantum computing technologies and experiment with different configurations. Collaboration between different groups, including physicists, computer scientists, and engineers, would then facilitate the development of co-design and integration processes.

However, *to develop quantum software compilation and stack solutions for quantum computing technologies other than superconducting qubits, it is important to understand the specific properties of each type of technology.* Next, we highlight some of these different properties by discussing the hardware characteristics of three prominent quantum computing technologies: (I) photonic qubits, (II) trapped ion qubits, and (III) neutral atom qubits.

(I) Photonic Qubits. Photonic qubits use photons as the qubits (known as qumodes in the context of photonic quantum computing). Photons are particles of light that have the unique property of being able to exist in a superposition of states. The properties of photonic qubits that are important to consider for quantum software compilation solutions are as follows [2], [8].

- **Output Usage Flexibility.** Because the output of a photonic qumode execution is distributed over an infinite basis, it gives greater flexibility in terms of selecting the best qumode states to optimize over, especially for algorithms based on Gaussian Boson Sampling (GBS).
- **High Interactivity.** Photonic qubits can be easily manipulated and transported over long distances, which makes them highly connected compared to other qubit types. This property is useful for developing quantum communication protocols and for implementing quantum error correction and mitigation techniques.
- **Long Coherence Times.** Photonic qubits have relatively long coherence times, which means they can maintain their quantum states for longer periods of time. This property is important for developing quantum algorithms that require long coherence times, such as quantum simulations and quantum cryptography.

(II) Trapped Ion Qubits. Trapped ion quantum computers use ions – that are trapped and controlled by electromagnetic fields – as the qubits. Unique properties of trapped ion qubits that are important to consider for quantum software compilation and stack solutions are as below [6], [9].

- **High Fidelity.** Trapped ion qubits have relatively high fidelity (as compared to superconducting qubits because ions are naturally quantum particles), which means that their quantum states can be manipulated and measured with higher accuracy, producing higher output fidelity.
- **Limited Connectivity.** Trapped ion qubits have limited connectivity, which means that the number of ions that can be connected and controlled is limited. This property is important to consider for developing quantum algorithms that require high connectivity, such as quantum simulations and quantum optimization problems.
- **Long Execution Times.** Trapped ion qubits have relatively longer coherence times, which means that their quantum states decay slowly; however, the operations take much longer to runs, resulting in longer execution times. This property makes it challenging to implement long quantum algorithms.

(III) Neutral Atom Qubits. Neutral atom qubits use neutral atoms (as opposed to ions) as the qubits. Neutral atom qubits have some important properties that must be considered for quantum software compilation and stack solutions [1], [5].

- **Multi-Qubit Gates.** Neutral atom quantum computers have the ability to implement gates involving more than two qubits, which is not possible with superconducting qubits. This opens up opportunity to run larger and fewer operations, but also creates challenges as operations need to be serialized to avoid blockades.
- **High Connectivity.** Neutral atom qubits can be connected and controlled using optical tweezers and magnetic fields, which makes them highly connected compared to other qubit types. This property reduces the number of required qubit state transference procedures, which can potentially improve the output fidelity.
- **Long Coherence Times.** Similar to photonic qubits, neutral atom qubits have relatively long coherence times, which means that they can maintain their quantum states for longer periods of time. This property is important for developing quantum algorithms that require long execution – and therefore, coherence – times.

In summary, developing quantum software compilation and stack solutions for photonic qubits, trapped ion qubits, and neutral atom qubits requires considering their unique properties. These properties include high connectivity, long coherence times, multi-qubit gates, high fidelity, and output usage flexibility. Developing new tools and techniques that can optimize for these properties is essential for advancing quantum computing technologies.

IV. Assessment

The success of developing quantum software compilation and stack solutions for other quantum computing technologies can be evaluated based on several metrics. For example, the primary metric to assess would be the fidelity of the output, which measures the impact of hardware noise effects on the ability to achieve correct and meaningful output. Other metrics such as the efficiency of the compiled code (number of gates and length of the critical path), the compilation overhead, and the speed of execution should also be used to evaluate the performance of different software solutions. Furthermore, the scalability of these solutions can be evaluated based on their ability to handle increasingly larger quantum algorithms.

V. Timeliness and Maturity

Developing software solutions for other quantum computing technologies is essential for the continued progress of the field. The timely advancements in the development of photonic qubits, trapped ion qubits, and neutral atom qubits have made it possible to build quantum computers with unique properties that can overcome some of the limitations of superconducting qubits. However, developing software solutions that can optimize for these properties requires the development of new compilation tools and techniques. The impact of success can be significant in various industries, such as cryptography, scientific computing, high-performance computing, finance, healthcare, and materials science.

In conclusion, developing quantum software compilation and stack solutions for other quantum computing technologies can enhance the scalability, efficiency, and performance of quantum computing. To achieve this goal, there is a need to develop new tools and techniques and create modular and portable software stack. The success of these solutions can be evaluated based on their ability to handle large and complex algorithms and their scalability to practical applications.

References

[1] Loïc Henriet, Lucas Beguin, Adrien Signoles, Thierry Lahaye, Antoine Browaeys, Georges-Olivier Reymond, and Christophe Jurczak. Quantum Computing with Neutral Atoms. *Quantum*, 4:327, 2020.

[2] Nathan Killoran, Josh Izaac, Nicolás Quesada, Ville Bergholm, Matthew Amy, and Christian Weedbrook. Strawberry Fields: A Software Platform for Photonic Quantum Computing. *Quantum*, 3:129, 2019.

[3] Prakash Murali, Norbert Matthias Linke, Margaret Martonosi, Ali Javadi Abhari, Nhung Hong Nguyen, and Cinthia Huerta Alderete. Full-Stack, Real-System Quantum Computer Studies: Architectural Comparisons and Design Insights. *arXiv preprint arXiv:1905.11349*, 2019.

[4] Tirthak Patel, Abhay Potharaju, Baolin Li, Rohan Roy, and Devesh Tiwari. Experimental Evaluation of NISQ Quantum Computers: Error Measurement, Characterization, and Implications. In *2020 SC20: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 636–650. IEEE Computer Society, 2020.

[5] Tirthak Patel, Daniel Silver, and Devesh Tiwari. GEYSER: A Compilation Framework for Quantum Computing with Neutral Atoms. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, pages 383–395, 2022.

[6] Ivan Pogorelov, Thomas Feldker, Ch D Marciniak, Lukas Postler, Georg Jacob, Oliver Krieglsteiner, Verena Podlesnic, Michael Meth, Vlad Negnevitsky, Martin Stadler, et al. Compact Ion-Trap Quantum Computing Demonstrator. *PRX Quantum*, 2(2):020343, 2021.

[7] John Preskill. Quantum Computing in the NISQ Era and Beyond. *Quantum*, 2:79, 2018.

[8] Shuntaro Takeda and Akira Furusawa. Toward Large-Scale Fault-Tolerant Universal Photonic Quantum Computing. *APL Photonics*, 4(6):060902, 2019.

[9] Xin-Chuan Wu, Dripto M Debroy, Yongshan Ding, Jonathan M Baker, Yuri Alexeev, Kenneth R Brown, and Frederic T Chong. Tilt: Achieving Higher Fidelity on a Trapped-Ion Linear-Tape Quantum Computing Architecture. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 153–166. IEEE, 2021.

A Quantum Approach for Efficient Biomedical Data Analysis

T. Perciano[1,*], J. Balewski[1] (balewski@lbl.gov), R. Van Beeumen[1] (rvanbeeumen@lbl.gov), E.W. Bethel[2] (ewbethel@sfsu.edu), D. Camps[1] (dcamps@lbl.gov), D. Huang[2] (danehuang@sfsu.edu)
[1]Computing Sciences Area, Lawrence Berkeley National Laboratory, [2]Computer Science Department, San Francisco State University
[*]Primary corresponding author: T. Perciano (tperciano@lbl.gov)

ASCR Basic Research Needs in Quantum Computing and Networking

Topic

Applications and Algorithms (**biomedical data analysis**): Innovative approaches leveraging the growing area of quantum information science (QIS) to develop quantum data representation and analysis methods tailored to various data types targeting the area of biomedical data analysis using quantum hardware and high-performance computing (HPC).

Challenges

Identifying the bottlenecks in analysis pipelines The first step when developing algorithms for biomedical data analysis is to identify specific bottlenecks in the data analysis pipeline. Biomedical data analysis often involves multiple steps, such as data preprocessing, normalization, feature extraction, modeling, etc. For each step, a specific algorithm can be used depending on the type of data and research questions. Identifying specific steps/algorithms for which classical computational complexity constrains problem size [1] and that may provide some quantum advantage to accelerate sciences is desired. For these specific algorithms, issues such as the rapid growth of data complexity and volume due to increase in sensor resolution/readback rates and the limits of classical approaches can be encountered. Finally, it is important to note that a key aspect of identifying bottlenecks in biomedical data analysis pipelines is the need of input from various stakeholders, including domain experts, data scientists, computational biologists, and quantum computing experts.

Complexity of biomedical data [2]
Heterogeneity. Biomedical data is inherently complex, coming from different sources including genomic and imaging data, neurological signals, and others. Each type of data has its own characteristics, processing requirements, and necessary analysis techniques.
Volume. Data can be extremely large, especially when dealing with high-throughput technologies such as genomics, transcriptomics, and medical imaging. Handling and processing such massive datasets efficiently is a challenge for quantum algorithms, as the number of qubits and the computational complexity grow rapidly with increasing data size. Also, quantum computers currently have limited memory capacity, which may not be sufficient to store these large amounts of data. On the other hand, for some biomedical applications the amount of data is rather small, which poses a different set of challenges, including, overfitting, limited statistical power, uncertainty, data sparsity, and data quality.
Data Preprocessing. Biomedical data often requires preprocessing steps to remove noise, correct bias, and transform data into a format that is suitable for further analysis. Quantum algorithms have to handle with these preprocessing steps while preserving the integrity and accuracy of the data.

Data representation [3, 4, 5, 6] Quantum computing employs a fundamentally different paradigm compared to classical computing, using qubits instead of bits. Developing effective data representation schemes that can encode biomedical data into quantum states is crucial for the development of quantum algorithms. Finding the optimal representation may be challenging, particularly for complex and diverse biomedical data types.

Data privacy As with any computational method in the biomedical domain, issues related to intellectual property and data privacy must be considered. Developing quantum algorithms that comply with data privacy regulations and protect sensitive patient information is crucial [7].

Mathematical foundations Mathematical concepts and foundations play a key role in quantum algorithms for biomedical data analysis. Quantum computing theory is based on a fundamentally different set of principles compared to classical computing. Consequently, understanding the mathematical concepts involved and developing new basic mathematical approaches becomes essential for new quantum algorithms. For example, translating biomedical problems into a suitable mathematical representation that can be effectively solved by quantum algorithms is challenging. Other challenges in this space involve the algorithm design, error analysis, efficient optimization techniques, among others.

Integration with existing pipelines Developing a quantum algorithm is only part of the solution; it must be seamlessly integrated with the existing biomedical data analysis pipeline. This requires collaboration among quantum computing experts, biomedical researchers, and data scientists to ensure compatibility and efficiency.

Hardware limitations [8] Current quantum computing hardware is still in its nascent stages, with limited qubits and relatively high error rates. As a result, developing quantum algorithms that can run efficiently on current hardware is a challenge.

Validation and benchmarking It is essential to validate the performance of new quantum algorithms against established classical methods. This requires well-defined benchmarking protocols and datasets to ensure the quantum algorithms are providing the expected improvements.

Workforce development and collaboration Quantum computing is a rapidly evolving field, and there is a significant skill gap between quantum computing experts and biomedical researchers. Bridging this gap will be crucial for the successful implementation of quantum algorithms in biomedical data analysis pipelines.

Opportunities

Quantum algorithms tailored to biomedical data analysis Development of novel algorithms and techniques leveraging the unique capabilities of quantum computing to address the challenges posed by complex and high-dimensional biomedical data [9]. The potential advantages can be related to exponential speedup and parallelism, to improve the efficiency, accuracy, and scalability of biomedical data analysis tasks. Key research opportunities include the development of quantum algorithms for feature extraction and selection, pattern search and matching, data preprocessing, and the integration of heterogeneous data types. Additionally, the exploration of quantum optimization techniques and quantum-inspired heuristics for solving complex biomedical optimization problems could lead to significant improvements in areas like drug discovery, protein folding, and biomarker identification.

Efficient quantum data encoding methods Development of new techniques for representing and manipulating diverse biomedical data types in the quantum computing domain. The aim is to address the challenges associated with data size, format, and complexity while leveraging the unique capabilities of quantum computing to enhance the efficiency and accuracy of biomedical data analysis tasks. Key research opportunities include the investigation of various quantum data encoding schemes to find the most suitable and efficient ways to encode different types of biomedical data. Quantum data compression techniques can be explored to reduce the size of biomedical datasets while preserving essential information, which would be beneficial for processing large-scale data with limited qubit resources. One key opportunity here is to create simple circuits suitable for NISQ devices.

Privacy-preserving data analysis Development of methods able to protect sensitive information in biomedical data, still enabling accurate and meaningful analysis, to address the challenges associated with data privacy, confidentiality, and regulatory compliance. Key research opportunities include the investigation of quantum algorithms and techniques that can ensure data privacy and comply with relevant regulations. Researchers could explore methods such as secure multi-party computation, differential privacy, and homomorphic encryption, which allow for the analysis of encrypted or obfuscated data without exposing the underlying sensitive information. Moreover, the development of privacy-preserving federated learning techniques, which enable the training of machine learning (ML) models on decentralized data sources, minimizing the need for data sharing and centralization. Researchers can also explore the development of secure and privacy-preserving quantum computing frameworks for biomedical data analysis.

Adaptive analysis pipelines Developing flexible and dynamic data analysis methods adaptable to the unique characteristics and requirements of diverse biomedical data types and tasks. The aim is to address the challenges associated with the heterogeneity and complexity of biomedical data and the need for efficient and scalable data analysis pipelines that can be easily adapted to different contexts and applications. Key research opportunities include the investigation of ML techniques for automating the identification and optimization of bottlenecks in biomedical data analysis pipelines. Researchers could explore the development of adaptive algorithms and methods that can dynamically adjust their parameters, features, or models based on the specific data at hand or the desired analysis goals. This may involve the integration of active learning, transfer learning, and meta-learning approaches into the biomedical data analysis pipelines. Furthermore, researchers can work on creating frameworks and tools that facilitate the seamless integration of quantum and classical algorithms, as well as the development of hybrid quantum-classical methods that can leverage the strengths of both computing paradigms.

Quantum machine learning Developing novel quantum algorithms to enhance the performance and efficiency of ML tasks in the context of biomedical data analysis [10]. The aim is to address the challenges associated with high-dimensional, noisy, and complex data and the need for efficient, scalable, and accurate ML methods tailored to specific biomedical applications. Key research opportunities include the investigation of quantum ML algorithms, such as quantum support vector machines, quantum neural networks, and quantum clustering methods, that can handle the complexity and noise inherent in biomedical data. Researchers could explore the development of quantum-inspired heuristics and optimization techniques for training ML models, which could lead to significant improvements in convergence speed and model accuracy. Moreover, the study of quantum-enhanced reinforcement learning and unsupervised learning techniques can enable the development of more advanced algorithms for pattern recognition, feature extraction, and knowledge discovery in biomedical data.

Workforce development Addressing the need for skilled professionals who can effectively apply quantum computing and advanced data analysis techniques to the biomedical domain. The aim is to bridge the gap between the growing demand for expertise in quantum computing, ML, and biomedical data analysis and the current availability of trained professionals who can tackle the challenges associated with complex biomedical data. Key research opportunities include the development of interdisciplinary training programs, workshops, and courses that bring together experts from fields such as quantum computing, ML, computer science, and biomedical research. These educational initiatives can help cultivate a new generation of researchers who are well-versed in both the theoretical and practical aspects of quantum computing and biomedical data analysis. Additionally, fostering collaborations and partnerships between academia, industry, National Labs, and healthcare organizations can create opportunities for knowledge exchange, internships, and real-world problem-solving experiences, which are crucial for building a skilled workforce.

Assessment

Evaluation would be based on comparing performance of the new methods against established benchmarks using well-known classical counterparts. Also, evaluation based on pre-defined criteria should also be taken into account, especially when considering specific biomedical problems. Success can be defined based on improved accuracy, efficiency, scalability, interpretability, robustness, or generalizability of the solutions when applied to real-world biomedical data analysis tasks. Standardized datasets, performance metrics, and validation techniques should be employed to assess the impact of the proposed solutions on addressing the specific challenges associated with biomedical data analysis. It is important to note that evaluating the new algorithms using specific quantum-related metrics is essential. For example, how NISQ-friendly is the algorithm? How well the algorithm can use the available quantum resources? Moreover, the adoption and integration of these solutions into existing biomedical data analysis pipelines and their ability to provide meaningful insights for researchers and clinicians can also serve as indicators of success. In general, the successful evaluation of potential solutions hinges on their capacity to advance the state-of-the-art in biomedical data analysis and facilitate a deeper understanding of complex biological processes and phenomena.

Timeliness, maturity, impact

The combination of advancements in QIS [11], biomedical data explosion, interdisciplinary collaboration, and the unique DOE multidisciplinary environment and expertise makes now the right time to pursue research in quantum algorithms for biomedical data analysis. The rapid progress in quantum computing and ML has opened up new possibilities for developing innovative algorithms for biomedical data analysis, enabling researchers to tackle previously intractable problems and develop more efficient and accurate algorithms. The growth of biomedical data from diverse sources has increased the need for advanced data analysis techniques that can handle the complexity, high-dimensionality, and noise inherent in these datasets. This data explosion creates a strong demand for novel algorithms and methods capable of extracting meaningful insights from vast amounts of data. Some biomedical areas show already the potential for quantum advantage (polynomial and super-polynomial) for Hamiltonian simulation, matrix inversion, unstructured search, dynamic programming and ML [9]. The growing interest in interdisciplinary collaboration between fields such as computer science, statistics, and biomedical research has facilitated the exchange of ideas, techniques, and expertise. This collaboration enables researchers to identify and address the challenges associated with biomedical data analysis more effectively, driving progress in this domain. The impact of success in this area will be far-reaching, with the potential to revolutionize healthcare, our understanding of complex biological processes, and broad applicability of the developed algorithms to other scientific fields relevant to DOE.

References

[1] S. Baichoo and C. A. Ouzounis, "Computational complexity of algorithms for sequence comparison, short-read assembly and genome alignment," *Biosystems*, vol. 156-157, pp. 72–85, 2017.

[2] S. Dash et al, "Big data in healthcare: management, analysis and future prospects," *Journal of Big Data*, vol. 6, p. 54, Jun 2019.

[3] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, "Quantum embeddings for machine learning," 2020.

[4] M. G. Amankwah, D. Camps, E. W. Bethel, R. Van Beeumen, and T. Perciano, "Quantum pixel representations and compression for N-dimensional images," *Scientific Reports*, vol. 12, p. 7712, May 2022.

[5] J. Balewski, M. G. Amankwah, R. V. Beeumen, E. W. Bethel, T. Perciano, and D. Camps, "Quantum-parallel vectorized data encodings and computations on trapped-ions and transmons qpus," 2023.

[6] D. Camps and R. V. Beeumen, "FABLE: Fast Approximate Quantum Circuits for Block-Encodings," in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, (Los Alamitos, CA, USA), pp. 104–113, IEEE Computer Society, sep 2022.

[7] W. M. Watkins, S. Y.-C. Chen, and S. Yoo, "Quantum machine learning with differential privacy," *Scientific Reports*, vol. 13, p. 2453, Feb 2023.

[8] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, Aug. 2018.

[9] B. A. Cordier, N. P. D. Sawaya, G. G. Guerreschi, and S. K. McWeeney, "Biology and medicine in the landscape of quantum advantages," *Journal of The Royal Society Interface*, vol. 19, no. 196, p. 20220541, 2022.

[10] L. Wei et al, "Quantum machine learning in medical image analysis: A survey," *Neurocomputing*, vol. 525, pp. 42–53, 2023.

[11] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press, 2010.

Long-term cybersecurity applications enabled by quantum networks

Nicholas A. Peters,[1,*] Muneer Alshowkan,[1] Joseph C. Chapman,[1] Raphael C. Pooser,[1] Nageswara S. V. Rao,[2] and Raymond T. Newell [3]

[1] *Quantum Information Science Section, Oak Ridge National Laboratory, Oak Ridge, TN 37831*
[2] *Advanced Computing Methods for Engineered Systems Section, Oak Ridge National Laboratory, Oak Ridge, TN 37831*
[3] *MPA-Quantum, Los Alamos National Laboratory, Los Alamos, NM 87545*

[*] *petersna@ornl.gov*

Topics. (Primary) applications and (Secondary) integration for quantum networks

Challenge. Fault-tolerant universal quantum computing has the potential to make the majority of current cryptographic protocols obsolete. This is colloquially known as "the quantum computing threat." As a result, crucial cryptographic functions that provide confidentiality, integrity, and availability of the communications which underpin global infrastructures are potentially at risk. This risk extends to the classical communications systems required to operate a quantum computer or quantum network; control signals and human-readable data are classical, and are therefore just as vulnerable as any other IT system to eavesdropping, spoofing, or other cyber attack. As we develop and deploy quantum networks, it will be essential to include security in their design from the beginning, rather than as a retrofit at the end.

There are two lines of defense against a cryptography-breaking quantum computer: classical cryptography systems which derive their strength from math problems which remain hard even with a quantum computer (often called Post-Quantum Cryptography, PQC); and quantum cryptographic systems (QCS) which derive their strength from the fundamental laws of physics. The first category, PQC, is under active development worldwide, including an ongoing competition sponsored by NIST to select and standardize quantum-safe cryptosystems [1]. On the other hand, QCS require development of separate hardware for their deployment; examples include quantum key distribution (QKD), which is the most mature QCS protocol, quantum digital signatures (QDS), and quantum secret sharing (QSS).

As a result of new hardware development for QCS, there are multiple research challenges that need to be addressed for QCS to realize its full potential. As QCS is a hardware-based solution, it is currently very expensive. For example, most discrete-variable QCS systems (e.g., encoding in polarization or time bin) utilize direct single photon detection (DD) with costly single-photon detectors (at least in the context of most modern telecommunications fiber networks). Additionally, DD-QCS can be severely limited by Raman scattering of classical light used to carry data [2]. Continuous variable (CV) approaches (encoding in amplitude and phase) utilize homodyne detection which is more cost effective, relatively immune to Raman scattering [3], and highly efficient during room-temperature operation. As a result, integration of DD-QCS into optical networks is challenging without very strict limitations on conventional data signals carried in the same fiber. In contrast, CV-QKD can be deployed with multiple optical channels carrying commercial levels of data. However, the DD-QCS is much more mature than CV-QCS, and for example, additional assumptions are frequently made for CV-QKD security about the detection process. In either case, most QCS systems are still expensive laboratory experiments or in bulky rack-mounted boxes with limited ruggedness for deployment. Moreover, it is an important open research question as to how to best securely implement and certify QCS, for example, so that side channels do not leak unintended information.

In addition, QCS assumes that an authenticated classical communications channel is available for the after-quantum transmission processing of the protocol. Much of the current cryptography infrastructure is public (asymmetric) key based whereas QKD delivers (private) symmetric keys, so QKD is not a drop-in replacement of current infrastructure. While QDS and QSS are multi-party protocols, they are not drop in replacements for existing cryptography methods either. As a result, how to authenticate the classical conventional channel and how to best utilize QCS in existing infrastructures remain a research challenge.

Opportunity. PQC is conventional cryptography thought to be secure against significant quantum computers and is hoped to be secure against foreseeable technology developments for at least several decades. On the other hand, QCS could potentially enable security for much longer time scales since the security is dependent on physics, which is technology agnostic, instead of on computational difficulty like PQC. This is an attractive feature for securing critical infrastructure, which could include science networks as well as energy systems, which are often expected to last for decades or longer.

Furthermore, QCS networks can benefit from using satellite-based quantum networks because satellites are difficult to access and can be monitored to ensure they remain physically secure, at least as secure as any ground station and probably more secure. The trade-off is that ground stations are required to communicate with the satellite; but, rugged portable ground stations are being commercialized. Satellite links could also enable longer-distance QCS links sooner than waiting for quantum repeaters to be developed [4], even benefiting from satellites in geostationary orbit that allow for more continuous key generation [5]. As a concrete example, the development of small rugged high-performance QKD satellites serving as “trusted nodes” would provide almost immediate practical benefit by distributing usable cryptographic keys to interested users on global scale.

Moreover, measurement-device-independent and device-independent implementations [6,7] significantly reduce the security requirements and assumptions. These more advanced protocols are designed to be secure even if certain parts of the hardware are not able to be located in a secure location. There have been demonstrations of measurement-device-independent QKD. Given the further increased security of these measurement-device-independent implementations over standard QCS (which already has advantages over other cryptography protocols relying on assumed computational difficulty) they could enable more flexibility in how QCS systems are deployed.

Assessment. In contrast to much of the rest of the world, in recent years, QCS research has not been a major focus in the US, where the focus is on PQC. The National Institute of Standards and Technology (NIST) is working to standardize PQC to counter the quantum computing threat [1]. It does not seem possible to standardize QCS and related technologies, using the same process as PQC as they are rooted in fundamentally different ideas. In addition, QCS is relatively immature compared to conventional cryptography. Nevertheless, for QKD in particular, there are natural metrics of secure key rate which are common to compare performance between different implementations. Additionally, device-independent and measurement-device-independent QCS protocols provide security verification that can be certified via loop-hole free Bell tests and Bell state measurements, respectively [6–8].

Timeliness or maturity. Despite the research challenges, the promise of long-term security independent of computational capability has caused QCS to be the focus of numerous academic research and corporate development programs globally. And while it is relatively immature compared to conventional cryptography, it is already a commercialized nearer-term application enabled by quantum networks. Even though there are several commercial offerings, much research and development must be done to close the gaps we describe. This research would hopefully make products more secure, enable their certification, enable longer communications distances, and increase their secure key rates. When fully mature, QCS protocols could enable long-term security of future quantum networking infrastructures.

References

1. L. Chen, S. Jordan, Y.-K. Liu, D. Moody, R. Peralta, R. Perlner, and D. Smith-Tone, “Report on post-quantum cryptography,” NIST Special Publ. **NISTIR 8105** (2016).
2. N. A. Peters, P. Toliver, T. E. Chapuran, R. J. Runser, S. R. McNown, C. G. Peterson, D. Rosenberg, N. Dallmann, R. J. Hughes, K. P. McCabe, J. E. Nordholt, and K. T. Tyagi, “Dense wavelength multiplexing of 1550-nm QKD with strong classical channels in reconfigurable networking environments,” New J. Phys. **11**, 045012 (2009).
3. B. Qi, W. Zhu, L. Qian, and H.-K. Lo, “Feasibility of quantum key distribution through a dense wavelength division multiplexing network,” New J. Phys. **12**, 103042 (2010).
4. Q. Zhang, F. Xu, Y.-A. Chen, C.-Z. Peng, and J.-W. Pan, “Large scale quantum key distribution: challenges and solutions,” Opt. Express **26**, 24260–24273 (2018).
5. J. C. Chapman, C. C. Lim, and P. G. Kwiat, “Hyperentangled time-bin and polarization quantum key distribution,” Phys. Rev. Appl. **18**, 044027 (2022).
6. H.-K. Lo, M. Curty, and B. Qi, “Measurement-device-independent quantum key distribution,” Phys. Rev. Lett. **108**, 130503 (2012).
7. A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani, “Device-independent security of quantum cryptography against collective attacks,” Phys. Rev. Lett. **98**, 230501 (2007).
8. P. Bierhorst, E. Knill, S. Glancy, Y. Zhang, A. Mink, S. Jordan, A. Rommal, Y.-K. Liu, B. Christensen, S. W. Nam, M. J. Stevens, and L. K. Shalm, “Experimentally generated randomness certified by the impossibility of superluminal signals,” Nature **556**, 223–226 (2018).

Measuring progress towards useful quantum computation using application-centric benchmarks

Timothy Proctor (tjproct@sandia.gov), Stefan Seritan, Robin Blume-Kohout, Andrew Baczewski

Quantum Performance Laboratory, Sandia National Laboratories

Topic: applications, algorithms, codesign and integration

Challenge: A quantum computer's information processing capabilities are limited by hardware errors (aka noise), and these errors are often complex and only partially understood.[1] This makes it challenging to understand each quantum computer's computational power, compare different processors, and both measure and guide technological progress. One way to quantify a processor's overall performance is to run a suite of test circuits—an approach known as *holistic benchmarking*.[1-11] Holistic benchmarks can be based on a wide variety of circuits, such as random[1] or algorithmic circuits,[2-8] and the power of a holistic benchmark is that it measures the impact of *all* of a processor's errors (known and unknown) on the benchmark's circuits. Holistic and application-centric benchmarking suites[2-4] have the potential to enable the quantification and comparison of quantum computing hardware's actual computational capabilities. However, some crucial challenges remain. Most existing application-centric benchmarks have not addressed important, foundational problems in the science of quantum computer benchmarking, and this severely hampers their utility.

An application-centric benchmarking suite's usefulness depends on *which* applications (i.e., algorithms and instances) it contains and *how* each application is turned into a benchmark. Importantly, an algorithm (or even an application) does not itself constitute a benchmark,[2] and—as we highlight below—standard approaches to turning an algorithm into a benchmark often result in impractical, unreliable, and opaque benchmarks. Here we highlight four important and urgent challenges to creating benchmarking suites that address DOE's needs:

1) *Benchmarking suites from curated DOE-relevant problems.* If an application-centric benchmarking suite is to measure progress towards useful quantum computation, the algorithms and problem instances on which it is based must be carefully curated to correspond to impactful applications. However, many existing benchmarking suites are based on a somewhat arbitrary selection of algorithms and problem instances, and so research to curate DOE-relevant applications for quantum computers will be essential if the community is to create benchmarking suites that serve DOE's needs.

2) *Benchmarks that scale to many-qubit processors.* The standard approach to creating a benchmark from an algorithm involves running circuits that implement that algorithm. But useful algorithms run circuits that cannot be classically simulated, resulting in a verification problem: how do you check the correctness of the output? Recent research has shown that this challenging problem *can* be solved[5-7]—by creating benchmarks from carefully altered versions of an algorithm's circuits—but many open questions remain.

3) *Benchmarks that are reliable measures of performance.* Existing application-centric benchmarking suites typically run an application for varied size problem instances (e.g., Shor's algorithm for n-bit integers with n varied), but a processor's performance on small instances of a problem has no rigorous relationship to progress towards solving large problem instances (and so classically intractable problems). Moreover, many common methods for quantifying a circuit's performance from data (e.g., classical fidelity) are surprisingly unreliable measures of the total error in that circuit. Benchmarking methods that reliably quantify a quantum computer's progress towards implementing a full algorithmic circuit are needed. Again, recent research has demonstrated that this is feasible,[5,8] but many open questions remain.

4) *Benchmarks that address the transition to fault tolerance.* Many quantum computer applications will require fault-tolerant operations enabled by quantum error correction (QEC), and an application implemented with near-term or full-scale QEC looks very different to its “bare” implementation with unprotected physical qubits. Benchmarking suites that ignore this distinction will not measure progress towards useful quantum computation. Benchmarking suites that include QEC routines will be part of the solution, but we conjecture that benchmarking methodologies that can meaningfully and transparently compare the performance of fault-tolerant and bare computations are also possible and necessary.

Opportunity: The research community can create benchmarking suites that meet the needs of DOE—and the broader scientific community—by collaborations between algorithms researchers, “QCVV” (quantum characterization, verification, and validation) scientists, software developers working on the quantum computing stack, and experimentalists. General benchmarking methodologies that address the challenges outlined above must be developed, and these methods must be backed by rigorous theory and experimental validation. Researchers in the QCVV and quantum algorithms communities have begun (and continue) to tackle the foundational problems in benchmarking methods; for example, techniques have recently been developed that can create scalable and provably robust benchmarks from any algorithm.[5-8] But existing methods have a range of technical limitations, and there is much more work to be done. Critically, principled benchmarking techniques must be implemented in software and integrated into application-centric benchmarking suites being developed by the community. This will require scientific collaborations between scientists in distinct subfields (e.g., QCVV and algorithms researchers), and synergistic software development.

Assessment: Success consists of practical and reliable application-centric benchmarking suites that can measure progress towards useful quantum computation. These benchmarking suites should be: (i) implemented in robust software; (ii) based on a carefully curated sets of problems and applications for quantum computing that are relevant to DOE (and the broader scientific community); and (iii) use cutting-edge benchmarking methodologies that are backed by rigorous theory and tested in experiments.

Timeliness or maturity: Holistic, application-centric benchmarks will continue to proliferate over the next few years—whether or not they are built on rigorous methods and theory—as will their adoption by experimental groups and quantum computer users. The first, path-finding application-centric suites[2-4] have played an invaluable role in kick-starting this development, but the limitations in the foundations of existing suites must be corrected before they become de-facto standards. Moreover, recent scientific breakthroughs and experience make this possible now. There are initial solutions to many of the foundational scientific challenges to application-centric benchmarking,[1,5-8] pioneering application-centric benchmarking software has revealed (and addressed) many of the practical challenges to implementing benchmarks in general-purpose code,[2-4] and engagement from experimental teams is beginning to take off.[9-10] The community is now ideally positioned to begin integrating cutting-edge benchmarking methods with practical benchmarking suites, and to use lessons learned and novel scientific insights gained along the way to further improve the foundations of quantum computer benchmarking science.

References

1. T. Proctor et al., Measuring the Capabilities of Quantum Computers, Nat. Phys. 18, 75 (2021).
2. T. Lubinski el at., Application-Oriented Performance Benchmarks for Quantum Computing, IEEE Transactions on Quantum Engineering 4, 1 (2023).
3. T. Tomesh et al., SupermarQ: A Scalable Quantum Benchmark Suite, http://arxiv.org/abs/2202.11045.
4. A. Li and S. Krishnamoorthy, QASMBench: A Low-Level QASM Benchmark Suite for NISQ Evaluation and Simulation, http://arxiv.org/abs/2005.13018.
5. T. Proctor et al., Establishing Trust in Quantum Computations, http://arxiv.org/abs/2204.07568.
6. S. Ferracin et al., Experimental Accreditation of Outputs of Noisy Quantum Computers, Phys. Rev. A 104, 042603 (2021).
7. J. Hines and T. Proctor, Scalable Full-Stack Benchmarks for Quantum Computers, in preparation.
8. S. Seritan et al., Benchmarking Hamiltonian Simulation via Circuit Mirroring, in APS March Meeting 2022, Vol. 67, Number 3 (American Physical Society, 2022).
9. IonQ, https://ionq.com/resources/applications-as-performance-benchmarks-for-quantum-computing
10. M. Amico et al., Defining Standard Strategies for Quantum Benchmarks, http://arxiv.org/abs/2303.02108.
11. R. Blume-Kohout and K. C. Young, A Volumetric Framework for Quantum Computer Benchmarks, Quantum 4, 362 (2020).

Bring the Wisdom in Designing Internet Protocols to Quantum Networking

Chen Qian
University of California, Santa Cruz
qian@ucsc.edu

Shengwang Du
University of Texas at Dallas
dusw@utdallas.edu

Katia Obraczka
University of California, Santa Cruz
obraczka@ucsc.edu

Ruilin Zhou
University of California, Santa Cruz
rzhou39@ucsc.edu

Yuhang Gan
University of California, Santa Cruz
ygan11@ucsc.edu

Xiaoxue Zhang
University of California, Santa Cruz
xzhan330@ucsc.edu

***Abstract*—Quantum networking is an emerging field that attracts increasing attention from both the quantum physics and computer networking communities. In this paper we share our research experience on developing quantum entanglement routing protocols from the perspective of network engineers. We believe many design considerations and experience of the classic Internet can be brought to quantum networks and help to develop practical, efficient, and scalable quantum networks in the future.**

I. Topic

Recently, quantum networks have been experiencing rapid development as a new type of network architecture that uses special hardware (quantum repeater equipped with quantum memory) to enable the transmission of quantum bits (qubits) [1]. There are two main types of quantum networks. The first type transmits qubits in a hop-by-hop manner [2], similar to the "store-and-forward" packet switching networks. However, it requires all repeaters in the network to be trusted, which is only applicable for private networks and is an impractical assumption in future large-scale quantum Internet [3]. The second type relies on entanglement routing [4] [3] based on the DLCZ protocol [5], which attempts to first establish entanglements (called external links) between every pair of consecutive repeaters along the way from the source to destination, and uses these external links to establish an end-to-end source-destination (S-D) connection through *entanglement swapping* [6]. Fig. 1 shows an example of entanglement swapping. From a networking perspective [3], we can consider there are two *external links* (p, q) and (m, n) and one internal link (q, m) that connect the whole end-to-end path from S to D.

We have conducted research in the past three years in the entanglement routing problem [4] [3], which is a method to find end-to-end paths consisting of external and internal links. We suggest that entanglement routing can be considered on the network layer of a quantum network.

II. Challenges

We realize that many existing protocols proposed for entanglement routing are with oversimplified network models and impractical assumptions. In fact, many existing ideas and experience from classic networks and be used to improve the designs of quantum network protocol. We share a few projects conducted by us in the past three years.

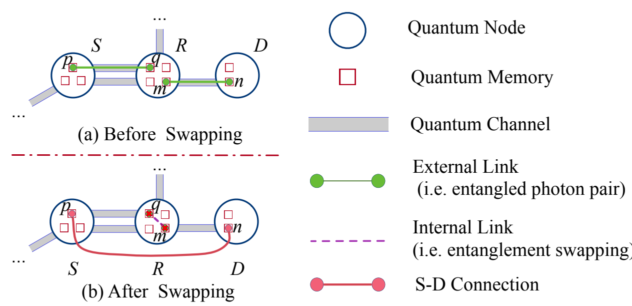

Fig. 1: Entanglement Swapping

III. Opportunity

Our first contribution to the problem was published to *ACM SIGCOMM* 2020 [3]. This work is the first work of a comprehensive protocol design specifically for entanglement routing in quantum networks, with new models, new metrics, and new algorithms, working on arbitrary network topologies. We present a comprehensive entanglement routing model that reflects the difference between quantum networks and classical networks and propose new entanglement routing designs that utilize the unique properties of quantum networks. We propose a few routing metrics that particularly fit quantum networks instead of using hop-count and physical distance. The proposed algorithms include realistic protocol-design considerations such as arbitrary network topologies, multiple concurrent sources and destinations to compete for resources, link state exchanges, and limited qubit capacity of each node, most of which have not been considered by prior studies. All these considerations are standard for routing protocol designs of classic network but was missing from quantum network before our work. Evaluation results show that the proposed algorithm Q-CAST increases the number of successful long-distance entanglements by a big margin compared to other known methods. A simulator with algorithm implementation, topology generation, statistics, and network visualization functions is available on this link [7]. More importantly, this study has encouraged more network researchers to study the entanglement routing problem and received more than 50 citations in less than three years. We present and clarify the models and problems of entanglement routing, with the comparison of similar terms and concepts used in classical network research.

Our second work focuses on the routing model. Most solutions of entanglement routing use the ***Syn**chronized **S**ingle-**t**ime-**s**lot model* (SynSts) [3]. The model strongly implies that all repeaters in the network carry similar hardware; hence their entanglements have similar lifetimes within a time-slot duration. We argue that this model does not align with the evolution of future quantum networks. Similar to the history of the Internet, building a large-scale quantum network for public users is a long-term task and consists of numerous *incremental deployments* of repeaters with heterogeneous hardware. The recent development of advanced quantum network hardware enables quantum memory to store photons with high fidelity for several minutes or even longer [8] [9].Therefore, the restriction that "the network must clear all external links at the end of each time slot" of the existing SynSts routing model will definitely result in sub-optimal results. We develop a quantum network routing framework with a new ***Syn**chronous **M**ulti-**t**ime-**s**lots (**SynMts**)* model to keep external links with heterogeneous time duration. In the new framework, we design a request management algorithm, `ReqUp`, which improves the network resources utilization in the network and a predictable links scheduling algorithm to manage the links in the network. With these two management algorithms, we show that many existing routing algorithms proposed for the SynSts model can be easily extended to run in the SynMts model.

The third project argues that all existing methods focus on an optimization goal of maximizing the routing throughput, i.e., the number of qubits delivered in a time unit. However, maximizing throughput is not sufficient to maximize the satisfaction of user applications, which is the ultimate goal of building a quantum network. Similar to the quality of service (QoS) requirements in classic Internet, the users of quantum networks also have requirements on the quality of the qubits they want to deliver. For example, quantum key distribution requires a certain level of fidelity of the delivered qubits to generate secret keys. Distributed quantum computing and synchronization require qubits to be delivered before a deadline. From the application perspective, we should focus on the *goodput*, which is defined as the number of qubits that are useful for applications in a time unit. QoS solutions in classical networks are different from quantum networks in that, metrics in classical networks like delay or bandwidth are either concave or additive. However, metrics in quantum networks like throughput and fidelity are not additive nor linear which requires a more complex design in routing algorithms. We design a QoS routing framework including a QoS routing process and scheduling process that can meet heterogeneous QoS requirements of the requests from multiple users concurrently. To our knowledge, this is the first entanglement routing solution with QoS metrics.

All three projects are motivated by the technologies in classic networks: network metrics and time slotted model, incremental deployment, and QoS requirements. We believe there are many opportunities of bringing experience and ideas in classic networks to quantum networks. We would like to discuss with the other attendees of ASCR Workshop on Quantum Computing and Networking. These opportunities are not limited to the routing problem and can span across the application layer, transport layer, network layer, and link layer, such as peer-to-peer communication, congestion control, routing, and link scheduling.

IV. Assessment

We have evaluated our projects on a customized discrete-event simulator [7]. We also investigated existing quantum network simulators and compare their features and advantages. We believe the quantum networking community should be able to offer an open-sourced, easy-to-use, and generalized simulator with available benchmarks of existing protocols. We would also like to discuss with other attendee about building a unified quantum network simulator just like NS3 for classic networks.

V. Timeliness

Quantum networking is an emerging field that attracts increasing attention from both the quantum physics and computer networking communities. At this time there are still many open problems in quantum networks, some of which might benefit from the experience and ideas of developing the classic networks. Hence we would like to encourage more computer networking researchers to participate in quantum network research using their demon expertise and want to discuss in the ASCR workshop on how to achieve it.

References

[1] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller, "Quantum repeaters: the role of imperfect local operations in quantum communication," *Physical Review Letters*, 1998.

[2] C. Elliot, "Building the quantum network," *New Journal of Physics*, 2002.

[3] S. Shi and C. Qian, "Concurrent entanglement routing for quantum networks: Model and designs," in *Proceedings of ACM SIGCOMM*, 2020.

[4] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha, "Routing entanglement in the quantum internet," *npj Quantum Information*, 2019.

[5] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics," *Nature*, 2001.

[6] R. Van Meter, T. Satoh, T. D. Ladd, W. J. Munro, and K. Nemoto, "Path selection for quantum repeater networks," *Networking Science*, 2013.

[7] "Source code of the quantum routing simulations." https://github.com/QianLabUCSC/QuantumRouting, 2019.

[8] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J.-N. Zhang, L.-M. Duan, D. Yum, and K. Kim, "Single-qubit quantum memory exceeding ten-minute coherence time," *Nature Photonics*, 2017.

[9] H. Bartling, M. Abobeih, B. Pingault, M. Degen, S. Loenen, C. Bradley, J. Randall, M. Markham, D. Twitchen, and T. Taminiau, "Entanglement of spin-pair qubits with intrinsic dephasing times exceeding a minute," *Physical Review X*, 2022.

Title: Quantum Computing and Networking: Embracing Clustered Quantum Computing (CQC) Systems for Tackling Qubit Limitations

a position paper submitted to DOE ASCR Workshop on Quantum Computing and Networking by

C. Qiao, R. Sutor, and S. Du

Topics: Quantum Computing Systems and Quantum Data Networking

Introduction: Quantum computing and networking are rapidly evolving research areas, poised to revolutionize the landscape of quantum information science and engineering (QISE) technology. As we witness the rapid progress in the development and deployment of quantum technologies, it is crucial for the scientific community to share a collective understanding of the challenges and opportunities that lie ahead. This position statement aims to provide an overview of the key challenges posed by the limited number of qubits in today's monolithic quantum computers (QCs) and the opportunities that clustered (or distributed) quantum computing (or CQC) systems offer. The goal is to foster discussion and promote research directions that can lead to breakthroughs in the field.

Challenge: The limited qubit count in today's monolithic QCs is a significant challenge, as most practical quantum applications require several orders of magnitude more qubits than what can be handled by these systems. Despite significant strides in increasing qubit counts and coherence times, monolithic QCs are still far from achieving the scale required for solving complex, real-world problems. For example, Xanadu has launched a QC called Borealis with 216 qubits, while IBM has built a 433-qubit QC chip called Osprey and plans to build a 1,121-qubit chip in 2023. However, real-world high-performance computing applications involving Quantum Annealing, Quantum Approximate Optimization, and Quantum Machine Learning may require hundreds of thousands of qubits or more. It is hard, if not impossible, to scale up the number of qubits significantly on a monolithic QC.

Opportunities: Clustered (or distributed) quantum computing (CQC) systems, consisting of multiple quantum processing units (QPUs) networked (or interconnected) together, offer a promising solution to overcome the qubit limitations in each monolithic QC. By harnessing the power of multiple networked QCs, such CQC systems can achieve a much higher qubit count and computational capacity, enabling the execution of complex quantum algorithms and applications. Even if we consider today's groupings of qubits as "quantum cores", we will still need to interconnect these cores to build large QPUs. In many cases, the concerns and techniques for the interconnects are the same for large scale quantum networking, such as transduction to photonic qubits.

To realize the full potential of CQC systems, significant advancements in quantum data networks (QDNs), which differ from quantum key distribution (QKD) networks and the Quantum Internet, are necessary. QDNs can enable transmission (or sharing) of quantum state information between different QPUs, allowing them to work together to perform complex computations. Research into efficient QDN protocols, entanglement distribution, and network architectures will be crucial in building an efficient and high-performance CQC system.

It is also essential to develop scalable and modular QPU technologies to support seamless integration of the QPUs into a CQC system. This will require advances in qubit technologies, quantum hardware design, fabrication techniques, and system integration approaches. Moreover, efficient management of quantum resources such as entangled photon sources, qubit memory, quantum switches and repeaters in QDNs, and effective quantum computing task partitioning, mapping and scheduling to achieve load balancing in the CQC systems must be developed. As quantum sensors become more pervasive over the next several years, the techniques developed for networking quantum computers will be used for distributing the data from the sensors for storage or processing.

Assessment: The most sensible approach to evaluating a CQC system is to use a high-fidelity simulator that supports not only physical layer characteristics but also QDN protocols. One of the main performance metrics is the size of the Hilbert space in a CQC system, formed by the number of qubits that can be effectively entangled with each other to perform meaningful quantum computing operations involving two or more qubits. We expect that a properly designed CQC with two QPUs, each having x and y qubits, and supporting Hilbert spaces of dimensions at most 2^x and 2^y, respectively, would increase the Hilbert space dimensions to the order of $2^x \times 2^y$, which would be a multiplicative increase, instead of a linear increase, when compared to the Hilbert space of each QPU alone.

Timeliness or maturity: Despite advances in various qubit technologies and related quantum hardware technologies, it becomes increasingly difficult to *scale up* the number of qubits in monolithic QCs. Recent advances in quantum interconnects, QDN protocols, photon-atom interaction, and quantum photon-atom modules provide opportunities to *scale out* the number of qubits by building a CQC. Moreover, more and more researchers with backgrounds in computer systems (including HPC and distributed computing systems) and computer networking are joining the QISE community to collaborate with researchers with backgrounds in quantum mechanics and physics, and other fields. We expect such collaboration to lead to new approaches and architectures for high-performance quantum computing systems.

Conclusion: The limitations posed by the qubit count in ’current monolithic QCs present significant challenges for the field of quantum computing. However, by embracing clustered (or distributed) quantum computing systems, the scientific community can overcome these limitations and unlock the true potential of quantum technologies. This approach requires advances in quantum data networking, modular QPU technologies, and system integration. It fostersa shared understanding of the problem space and will stimulate discussion that can drive progress in the years to come. Researchers with complementary expertise from academic, industry and government labs, funding agencies, and policymakers need to work together to capitalize on these opportunities and overcome the challenges. This will accelerate the development and adoption of quantum technologies and enable breakthroughs for multiple use cases in science, government, and industry.

Contact:

Chunming Qiao, SUNY Distinguished Professor, Computer Science and Engineering, University at Buffalo, State University of New York. Email: qiao@buffalo.edu

Robert S, Sutor, Adjunct Associate Professor, Computer Science and Engineering, University at Buffalo, State University of New York. Email: rsutor@buffalo.edu. Also affiliated with Infleqtion (rbob.sutor@infleqtion.com)

Shengwang Du, Francis S. and Maurine G. Johnson Chair Professor, Department of Physics, University of Texas at Dallas, Email: dusw@utdallas.edu

Growing the Science of Quality Assurance for Quantum Software Stacks

Akond Rahman
Department of Computer Science and Software Engineering (CSSE)
Auburn University, Auburn, AL, USA
akond.rahman.buet@gmail.com

***Index Terms*—compilation, error correction and mitigation**

I. Challenges

The goal of this document is to help practitioners in developing robust quantum programs by providing an overview of challenges and opportunities related to quality assurance for quantum software stacks.

We accomplish this goal by identifying three technical challenge areas related to quantum software stacks. We believe these overcoming these challenges are critical to achieve robust quantum software development. We describe these challenges as follows:

A. Challenge-1: Lack of effective program analysis tools to identify latent bugs in quantum software stacks

Latent defects are challenging to identify due to the unique nature of quantum software. Without identification and mitigation of such defects, quantum-based software remain susceptible to generating large-scale consequences. The complexity and size of quantum software stacks pose a significant challenge in detecting these latent defects. One approach to overcome this challenge could be manual code review, which requires significant expertise that is typically unavailable in the software engineering domain. Therefore, we need to develop effective program analysis tools that can identify latent defects in quantum software stacks.

B. Challenge-2: Stack-agnostic robust and performant quantum program generation

The use of quantum software stacks can prohibit developer productivity. According to De Stefano et al. [2] "*adapting ideal quantum circuits to available device architectures is a tough challenge for developers*". We believe there are multiple dimensions to this challenge, which are described as follows:

- *Code generation:* Accurate code generation that actually runs on quantum circuits.
- *Developer productivity:* Developer productivity may vary from one quantum stack to another. For example, a developer who is proficient in Python may feel more productive for the Python-based PyQuil [8] and Qiskit [9] quantum software stacks. However, a developer who is proficient in a non-Python programming languages, such as C/C++ may prefer a different quantum software stack.
- *Performance:* Performance-aware quantum software development is also challenging as quantum software stacks use emulators that are resource hungry [2]. Therefore, while addressing developer productivity, future research should consider development of quantum computer programs that are performant.

Therefore, to address this challenge the proposed technique must account for stack agnostism and performance so that developed research techniques can work for all popular quantum software stacks in an efficient fashion.

C. Challenge-3: Lack of foundational understanding of misconfigurations unique to quantum software stacks

Recent empirical research [7] shows that misconfigurations is a common defect category for quantum computer programs. However, there is a lack of understanding on the nature of misconfigurations for quantum software stacks. We hypothesize that the configuration space for quantum software stack is large enough that poses significant challenges to understand quantum-related misconfigurations. While existing research has addressed bug detection for quantum computer programs further research is needed to gain a foundational understanding of misconfigurations for quantum software stacks.

II. Opportunities

The above-mentioned challenges provide opportunities that we describe as follows:

A. Opportunity-1: Novel defect detection tools for quantum software engineering

The challenge related to latent defect provides the opportunity to develop novel tools that can identify latent defects. One possible approach is to use model checking. Model checking is a technique to check if a model represented as finite state machines of a computer program or a system meets a given specification [1]. The possible sequence of steps to apply model checking could be: (i) creation of a property knowledge base for quantum programs, (ii) gain understanding to derive reference models, (iii) generate counterexamples with model checking, and (iv) prune generated counterexamples.

B. Opportunity-2: Robust and performant program generation to improve quantum-based developer productivity

The challenge related to developer productivity provides the opportunity to use large language models (LLMs) for quantum computer program generations. As part of this activity, researchers can investigate the effectiveness of LLMs, such as ChatGPT [6] and CodeBERT [3]. Next, researchers can use techniques, such as prompt engineering [5] with code regeneration to generate performant quantum computer programs. For stack-agnosticism, researchers can generate intermediate representations so that generated code can work all popular quantum software stacks.

C. Opportunity-3: Pro-active detection of misconfigurations in the quantum software stack to avoid large-scale consequences

By understanding the nature of quantum-related misconfigurations we can develop novel techniques that can pro-actively detect misconfigurations in quantum software stacks. One opportunity is related to using configuration testing with combinatorial testing, where researchers can use pairwise testing to discover misconfigurations in quantum computer programs.

Realization of these opportunities will generate potentially transformative research for the domain of quantum computing by applying, evaluating, and re-designing established techniques in the field of formal methods, natural language processing, and combinatorics.

III. Assessment

We propose the following assessment plan:

- Assessment of the proposed model checking-based approach:
  - latent defect localization accuracy with metrics, such as precision, recall, and F-measure;
  - latent defect discovery rate, i.e., the proportion of counterexamples that actually lead to latent defects; and
  - counterexample generation time, which measures the time to generate counterexamples.
- Assessment of robust program generation techniques:
  - the proportion of generated programs that quantum software stacks can compile;
  - the proportion of lines, branches, and functions in quantum software stacks are covered using LLM-based program generation techniques; and
  - the amount of time it takes for software developers to complete quantum computing tasks with and without the assistance of LLM-based program generation techniques.
- Assessment of the proposed configuration testing techniques:
  - defect discovery rate that measures the defect finding ability of a testing technique;
  - test run rate that measures the time incurred by the testing technique;
  - defect find rate that is computed by dividing the total number of defects by total test hours;
  - passed test case rate that measures the number of passed tests per test execution; and
  - critical defects rate that measures how many of the defects identified by the testing technique are critical defects.

IV. Timeliness

The above-mentioned challenges and corresponding opportunities are timely as they all contribute to a resilient quantum software ecosystem. By addressing all of the challenges mentioned in Section I the quantum software stack can be more usable for developers. All of this will contribute to the nation's ongoing efforts in addressing the quantum needs. For example, as described in Section II-C, if the opportunities related to defect and misconfiguration detection is realized, then generated research will directly contribute to mitigating quantum errors, and area highlighted in the 2018 National Strategic Overview for Quantum Information Science [4]. Furthermore, realization of the second opportunity related to quantum-related developer productivity described in Section II-B will help in creating a quantum-ready workforce, and area also highlighted in the 2018 National Strategic Overview for Quantum Information Science [4].

References

[1] E. M. Clarke, T. A. Henzinger, H. Veith, R. Bloem *et al.*, *Handbook of model checking*. Springer, 2018, vol. 10.

[2] M. De Stefano, F. Pecorelli, D. Di Nucci, F. Palomba, and A. De Lucia, "Software engineering for quantum programming: How far are we?" *Journal of Systems and Software*, vol. 190, p. 111326, 2022.

[3] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, "CodeBERT: A pre-trained model for programming and natural languages," in *Findings of the Association for Computational Linguistics: EMNLP 2020*. Online: Association for Computational Linguistics, Nov. 2020, pp. 1536–1547. [Online]. Available: https://aclanthology.org/2020.findings-emnlp.139

[4] U. W. House, "National Strategic Overview for Quantum Information Science," https://www.quantum.gov/wp-content/uploads/2020/10/2018_NSTC_National_Strategic_Overview_QIS.pdf, 2018, [Online; accessed 19-April-2023].

[5] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, "Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing," *ACM Comput. Surv.*, vol. 55, no. 9, jan 2023. [Online]. Available: https://doi.org/10.1145/3560815

[6] OpenAI, "Introducing ChatGPT," https://openai.com/blog/chatgpt, 2023, [Online; accessed 19-April-2023].

[7] M. Paltenghi and M. Pradel, "Bugs in quantum computing platforms: an empirical study," *Proceedings of the ACM on Programming Languages*, vol. 6, no. OOPSLA1, pp. 1–27, 2022.

[8] R. S. Smith, M. J. Curtis, and W. J. Zeng, "A practical quantum instruction set architecture," *arXiv preprint arXiv:1608.03355*, 2016.

[9] R. Wille, R. Van Meter, and Y. Naveh, "Ibm's qiskit tool chain: Working with and developing for real quantum computers," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 1234–1240.

Quantum Computing Requires Application-level Programming Models that Deliver Performance

Raouf Dridi
Quantum Computing Inc.
rdridi@quantumcomputinginc.com

Pranav Gokhale
Infleqtion
pranav.gokhale@infleqtion.com

Dor Israeli
Quantum Machines
dor@quantum-machines.co

Tim Mattson
Intel
timothy.g.mattson@intel.com

Tom Lubinski[§]
Quantum Circuits
tlubinski@quantumcircuits.com

Santiago Núñez-Corrales[§]
NCSA/IQUIST, UIUC
nunezco2@illinois.edu

Steve Reinhardt[§]
Transform Computing
spreinhardt.tech@gmail.com

Yuval Sanders
University of Technology, Sydney
Yuval.Sanders@uts.edu.au

Michał Stechły
PsiQuantum
mstechly@psiquantum.com

Tobias Stollenwerk
Forschungszentrum Jülich
to.stollenwerk@fz-juelich.de

A. *Topic*

Our topic is the need for better application-level **programming models** developed through a hardware/software **codesign** process.

B. *Challenge*

The goal of the quantum computing industry is to deliver quantum computers (QCs) that are profitable and deliver performance for high-value problems vastly superior to that from classical computers. This can only happen via a vibrant ecosystem of applications written by programmers with backgrounds in application domains rather than specialists in the physics of quantum computing.

A programming environment implements a layered stack of models that serve as a bridge between the developer and the hardware. Given that application software has a lifespan exceeding that of any particular computer, this stack must work for all relevant hardware. There is no single way to define these layers, but we think in terms of the following:

- **Programming model**: Abstractions that help application developers understand algorithms and how they map onto software.
- **Execution model**: Abstractions of how software executes on idealized machines.
- **Hardware model**: Abstractions that map idealized machines onto actual systems.

These abstractions will emerge from a codesign process based on a deep understanding of how to derive the best performance from QCs and an understanding of the fundamental algorithms driving applications software. We will need to work *down* from the application level and the requirements of application developers, while we simultaneously work *up* from the *quantum hardware layer*, capturing what will deliver the best performance.

In our view, insufficient attention has been paid to understanding these fundamental abstractions. The QC community needs better application-level programming models and supporting abstractions, enabling computational scientists to access immensely capable QCs via mature software stacks. We expect that creating productive programming languages to harness the power of those QCs requires a deliberate and coordinated effort. We advocate for this renewed focus to begin now.

The circuit model is widely used and has consumed most of the intellectual effort in quantum software. We see two serious shortcomings to the circuit model.

First, circuits *overspecify* solutions, requiring unnecessary details about computer architecture to be explicitly stated when what matters is the meaning of the transformations that capture the essence of an algorithm. This hinders the development of quantum computer applications because, as with classical algorithm development, much of the design process involves the specification and selection of data structures. The technical work to translate these structures into bit-level operations drastically slows down the workflow and is highly prone to hard-to-detect programmer error that impacts performance analysis. The burden of producing quantum circuits thus drastically slows down the design process for novel quantum algorithms and renders the resulting analysis of their value untrustworthy.

Secondly (and paradoxically), circuits *underspecify* details for mapping efficiently to hardware, given a circuit that solves a problem. Circuits suggest a degeneracy of implementation: a given set of quantum gates translates into multiple possible execution sequences, several circuit reduction strategies, and many general optimizations even within the confines of a single pulse-level interface. Hence, the circuit model lacks enough expressiveness to enable the necessary fine-tuned control of the hardware. This means developers cannot be

[§]We nominate Steve Reinhardt, Santiago Núñez-Corrales, and Tom Lubinski for workshop attendance, space permitting.

confident that a circuit-model implementation will achieve good performance for a problem that is intuitively well suited for QCs.

Recognized problems with the circuit model are not just hypothetical. In the last year, many researchers working with quantum benchmarks have shifted their attention to pulse-level programming [1], [2]. This enables exploiting additional degrees of freedom about how, for example, a CNOT gate is implemented in shape and timing of pulses. The growing number of QCs that are not natively gate-based, including measurement-based (PsiQuantum) and analog (Xanadu, QuEra, D-Wave) QCs, may also benefit from a higher-than-circuit-level model.

These ideas take yet another twist when quantum error correction (QEC) becomes practical, since the circuit cost model for QEC is dramatically different from today's noisy intermediate-scale quantum (NISQ) computers. In fact, the gates that are challenging in a NISQ context, e.g., CNOT, are relatively easy in QEC, while the gates that are easy in a NISQ context, such as $R_z(\theta)$, are challenging in QEC [3], [4]. In a possible future of *Very Large Scale Quantum Integration* (VLSQI) in which devices reach thousands of logical qubits supported by millions of physical qubits, both issues will only be exacerbated: problems will be too complex to represent as verifiably correct circuits, and (possibly simultaneously) too vaguely expressed to optimize in the context of a given QC.

C. Opportunity

The quantum computing community has an opportunity to change the foundation of our software ecosystem. We need to design quantum hardware around software. The sooner we have a concept of the high-level models that programmers will use, the sooner we'll build useful hardware. We believe that the circuit model is inadequate for computational scientists to pose problems (application layer) or for software tools to deliver full performance from QCs (hardware layer) as the number of quantum bits scales to VLSQI. **Hence, we seek to find better models, models that help computational scientists reason about solutions to real-world problems while allowing tools to map software onto QCs with high performance.**

Our view is that programming QCs must become a widely practiced skill for the QC industry to succeed. We compare to the history of GPU programming. CUDA was introduced in 2006. Since then, a hardware/software codesign process has driven innovation so today, over 15 years after its introduction, GPU programming is a skill expected of an HPC programmer (with 2 million registered CUDA users in 2021). It may take more than a decade for QC to match the widespread use of GPUs, but we believe that order of magnitude of users must be our design target.

While the innovations in CUDA have been driven by a single commercial entity, we envision an open, community-driven process for the far more complex problem of hardware-software codesign in quantum computing. The community, by working together now, has an opportunity to create the layered stack of abstractions that will foster the emergence of a QC software ecosystem. This would only occur if we can support a separation of responsibilities. Reasoning in terms of low level elements, such as circuits, should be the task of computer architects. Such separations of responsibilities are well known in digital computers, for which computer organization details are hidden from the view of application and even most compiler developers by instruction-set architectures.

For the application layer in QC, we are faced with an exciting intellectual challenge. Once again, reasoning from analogy with classical computing, computational scientists often think in terms of linear algebra distinct from any specific implementation; having devised an algorithm in terms of linear algebra, they can then readily implement it for classical computers, knowing that high-performance implementations of common linear-algebraic functions are widely available. We have an opportunity to create such a conceptual model for quantum programming and implement it while quantum hardware is maturing, hopefully guiding hardware developments so they will be capable of running such models effectively.

D. Assessment

Success for this research can be measured by three criteria.

- **Application diversity**: How many application domains are covered by software written for quantum computers?
- **The developer community**: Who are writing these applications? Biologists should write biology applications. Chemists should write chemistry applications. Programming models are a failure if only specially trained quantum computing experts write applications.
- **Performance**: All the performance available from a quantum computer must be exposed. High level models that can only generate low performance code are a failure.

Finally, these criteria must be met across all key commercial QCs. A research agenda should be ambitious so the goal should be a layered stack of models that support programming systems that deliver productivity, performance, and portability.

E. Timeliness

Numerous companies are using QCs for small proof-of-concept projects that anticipate future industrially relevant applications. These efforts have been hindered by the necessity of working at the circuit level. Circuit models may, with Herculean efforts, succeed for the first few demo-applications, but for QCs to reach their full potential and impact US national security and economic interests, we need hundreds (or thousands) of applications that solve a variety of problems. We need better software stacks that deliver performance on industrial QCs. Finding expressive high-level models that can extract high performance from QCs might take many years of hard work with close involvement from hardware, compiler, runtime system, and application developers. With the availability of practical, fault-tolerant QCs expected by the end of the decade, we need to redouble work on these application-level models now.

REFERENCES

[1] D. Meirom and S. H. Frankel, “Pansatz: Pulse-based ansatz for variational quantum algorithms,” *arXiv preprint arXiv:2212.12911*, 2022.

[2] L. Ella, L. Leandro, O. Wertheim, Y. Romach, R. Szmuk, Y. Knol, N. Ofek, I. Sivan, and Y. Cohen, “Quantum-classical processing and benchmarking at the pulse-level,” *arXiv preprint arXiv:2303.03816*, 2023.

[3] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Physical Review A*, vol. 86, no. 3, p. 032324, 2012.

[4] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, “Efficient z gates for quantum computing,” *Physical Review A*, vol. 96, no. 2, p. 022330, 2017.

Towards a More Principled Quantum Computing Ecosystem

Francisco Rios* and Kathleen Hamilton

Oak Ridge National Laboratory, Oak Ridge TN 37380, USA

Topic

Primary: Models. Secondary: Codesign and Integration.

Challenge

The capability of quantum computing to solve certain types of computational problems much more efficiently than its classical counterpart [5] has resulted in an impressive amount of theoretical and applied research. And even though large-scale fault-tolerant quantum devices are not yet a reality, promising applications of quantum computing to very diverse fields such as quantum simulation [6], machine learning [7], and natural language processing [1], among many others, have emerged. However, the ad-hoc tools and techniques of the different fields make it difficult to transfer results and insights from one to the other, thus hindering the development of a robust, integrated quantum computing ecosystem.

For example, when trying to prove the correctness of quantum programs, some of the main difficulties of such a task arise not only from the complexity of the quantum algorithms that these programs are supposed to implement but also from the fragility of the quantum information on which they are supposed to operate. Debugging a classical program by analyzing its state is a powerful technique in classical computing. However, it is ineffective in a quantum context since observing the state of a quantum program collapses the state. Thus, we are faced with the challenge of *developing more encompassing abstract frameworks in which quantum and classical information processing can be modeled and reasoned about simultaneously.*

Opportunity

The continuous discovery of new applications and insights in quantum computing requires improved theoretical foundations to tame the ever-increasing complexity of the field and provide a more uniform framework in which different application domains can interact seamlessly. Thus further research on the core logical, mathematical, and computational assumptions and techniques used in quantum computing is appropriate and necessary to push the field forward. Hence our community is presented with an opportunity to *perform foundational work that has the potential to impact all areas of quantum information processing significantly.*

Until today, the design and specification of programming languages and the analysis of programs written in them have mainly relied on classical and intuitionistic logics. Moreover, the expected arrival of scalable quantum computers in the not-so-distant future has propelled the development of quantum programming languages [4], for which classical and intuitionistic logics are insufficient. For these languages to be adequate for quantum computing, they must be able to manage quantum information as a non-duplicable resource to prevent violations of the no-cloning property. As a consequence, linear logic [2], being a resource-sensitive formal system, has emerged as the leading logical framework for specifying and studying quantum programming languages. However, as the expressiveness of these languages increases, so does the sophistication of the logical systems needed

*riosfr@ornl.gov

to reason about them and the categorical structures required to model them. Thus, further research in substructural logics, of which linear logic is an example, and corresponding categorical semantics is needed to *ensure the robustness of current and future quantum programming languages and the quantum software stack built with them.*

Other relevant recent applications of substructural logics and categorical semantics, in terms of typed grammars and compact closed categories, respectively, have materialized in quantum natural language processing (QNLP)–the application of quantum computing to natural language processing tasks. At the field's core is the Categorical Compositional Distributional (DisCoCat) model, a powerful abstract framework that uses category theory to unify the compositional nature of grammatical structures and the distributional theory of meaning of natural languages. The framework is quite flexible and has found applications as diverse as protein analysis, question answering, and even music composition, to name a few [3]. These applications further highlight the need for more foundational research on the logico-categorical structures that *enable* them.

Assessment

There are several ways to evaluate success in this context. One of them would be the continuous development of high-level, correct-by-construction quantum programming languages that are expressive enough to significantly facilitate the implementation of quantum algorithms from their descriptions, as found in the literature. Another measure of success would be expanding formal verification tools to the quantum computing ecosystem. However, the most relevant measure of success would be the construction of richer logics and mathematical models geared not only to support but to expand the languages, tools, and applications mentioned above.

Timeliness

The rapid development of quantum computing and applications has resulted in many ad hoc frameworks and approaches. While successful in their fields of application, transferring results and insights from one application to another can be quite challenging. Hence, a more principled, unifying, and formal approach to the area is needed now to reduce the current landscape of siloed applications and algorithms. This methodology would account not only for a more coherent and robust quantum computing ecosystem but also for a more profound knowledge of the power, limitations, and potential applications of quantum computing itself.

References

[1] Bob Coecke et al. "Foundations for near-term quantum natural language processing". In: *arXiv preprint arXiv:2012.03755* (2020).

[2] Jean-Yves Girard. "Linear logic: its syntax and semantics". In: *London Mathematical Society Lecture Note Series* (1995), pp. 1–42.

[3] Raffaele Guarasci, Giuseppe De Pietro, and Massimo Esposito. "Quantum natural language processing: Challenges and opportunities". In: *Applied Sciences* 12.11 (2022), p. 5651.

[4] Michele Mosca, Martin Roetteler, and Peter Selinger. "Quantum Programming Languages (Dagstuhl Seminar 18381)". In: *Dagstuhl Reports*. Vol. 8. 9. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2019.

[5] Michael A Nielsen and Isaac Chuang. *Quantum computation and quantum information*. 2002.

[6] James D Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. "Simulation of electronic structure Hamiltonians using quantum computers". In: *Molecular Physics* 109.5 (2011), pp. 735–750.

[7] Zhikuan Zhao et al. "Bayesian deep learning on a quantum computer". In: *Quantum Machine Intelligence* 1 (2019), pp. 41–51.

Title: ML-Assisted Noise-Aware Quantum Circuit Compilation

Author: Samah Mohamed Ahmed Saeed, City College of New York, City University of New York,
email: ssaeed@ccny.cuny.edu

Topic: Compilation, Error Correction, and Mitigation

Challenge: Significant research and industrial efforts are focused on developing noise-aware quantum circuit compilation techniques to boost the performance of quantum algorithms. The compilation stack includes gate decomposition, circuit optimization, and quantum circuit mapping. They target minimizing the gate count, the circuit depth, or maximizing the fidelity of the physical quantum circuit by considering different noise models (e.g. [1]). Noise-aware compilation techniques typically rely on error rates computed using randomized benchmarking due to its scalability [2]. However, they fail to accurately predict the impact of quantum hardware errors on the output fidelity when executing circuits of different structures. As a result, existing noise-aware quantum compilation techniques informed by these noise models may provide a limited improvement in the output state fidelity of the quantum circuit. Table 1 illustrates an example of five different physical implementations of a quantum circuit, which share the same initial qubit assignment, executed on IBM Q Jakarta during the same calibration window. It shows that optimizing the circuit depth (D), the number of CNOT gates ($\#CX$), or the Estimated Success Probability (ESP) based gate and readout error rates do not yield the highest output fidelity (FID).

Table 1: Quantum circuits executed on IBM Q Jakarta.

Cir.	FID	D	#CX	ESP
M_1	73.05	28	19	0.72
M_2	67.58	25	27	0.68
M_3	70.90	36	17	0.74
M_4	65.14	35	21	0.72
M_5	67.87	25	23	0.70

While error mitigation techniques such as noise tailoring via randomized compiling, which suppresses coherent errors [3, 4], and dynamical decoupling of idle qubits [5] have been shown to improve the output fidelity they still can benefit from higher-level noise-aware compilation techniques that take into account the hardware noise to minimize the interaction with the quantum hardware. For example, while dynamical decoupling protocols are used to suppress idling errors, recent research has shown that the performance of different DD sequence approaches on different quantum algorithms is still unknown [6, 7] motivating the need for an adaptive approach to selectively insert the DD sequences into a subset of qubits to improve the output state fidelity of the quantum circuit. There is a need to integrate accurate and scalable quantum circuit fidelity predictor and noise modeling into different levels of the compilation stack to guide principled approaches to designing reliable quantum circuits, rather than treating errors as an afterthought; thus, maximizing the potential of quantum computing systems.

Opportunity: Heuristic methods are based on manually-crafted decision rules, which can change depending on the physical features of the hardware. Recent research has shown that Machine Learning (ML) can unlock more opportunities in compiler optimization by replacing complicated heuristics with ML policies. Recent efforts have focused on building ML models, i.e. reliability models, to predict the quantum circuit output fidelity based on Deep Learning including these models that can preserve the circuit topology and therefore can outperform noisy quantum circuit simulation in terms of scalability while maintaining high accuracy [8–11]. As a result, new efforts to integrate these models into the compilation stack have been demonstrated [12–14]. These approaches not only reduce the design space exploration but also better predict the reliability of the circuit to achieve better fidelity. We argue that embedding the circuit topology information could enable the models to reason about code with higher accuracy. Finally, the exploration of quantum ML-based compilation approaches can also accelerate the training and inference of classical ML methods. To this end, extensive collaborations between ML/AI and quantum computing experts will result in significant progress in enabling different quantum computing applications.

Assessment: Several metrics can assess the ML-assisted noise-aware compiler including repro-

ducibility, output fidelity, scalability, cost-function-dependent metrics, entanglement entropy, and the run time of the compilation approaches. The reproducibility can explore the correlation between the quantum circuit output and the device characterization. The output fidelity measures the difference between the ideal and the noisy output. The scalability can be assessed by measuring the correlation between the model prediction accuracy and the size of the quantum circuits. Since quantum circuits that are designed to suppress quantum hardware noise may not necessarily provide the best application-specific success metric (e.g. barren plateaus in parametric quantum circuits), a cost-function-dependent metric is also required to evaluate them. Entanglement entropy can be used to study how entanglement can affect the optimization of quantum circuit performance.
Timeliness or maturity: This research direction addresses Quantum and AI Initiative Acts by advancing quantum computing using AI and AI using quantum computing, respectively, and by preparing a quantum and AI-smart workforce. It is also aligned with the current advancement in building intelligent compilers for classical computing to improve system performance and the AI transformation into a more accessible technology.

References

[1] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, "Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers," in *Proceedings of ASPLOS*, 2019, pp. 1015–1029.

[2] E. Magesan, J. M. Gambetta, and J. Emerson, "Characterizing quantum gates via randomized benchmarking," *Phys. Rev. A*, vol. 85, p. 042311, Apr 2012.

[3] J. J. Wallman and J. Emerson, "Noise tailoring for scalable quantum computation via randomized compiling," *Physical Review A*, vol. 94, no. 5, nov 2016.

[4] A. Hashim, R. K. Naik, A. Morvan, J.-L. Ville, B. Mitchell, J. M. Kreikebaum, M. Davis, E. Smith, C. Iancu, K. P. O'Brien, I. Hincks, J. J. Wallman, J. Emerson, and I. Siddiqi, "Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor," *Phys. Rev. X*, vol. 11, p. 041039, 2021.

[5] B. Pokharel, N. Anand, B. Fortman, and D. A. Lidar, "Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits," *Phys. Rev. Lett.*, vol. 121, p. 220502, 2018.

[6] P. Das, S. Tannu, S. Dangwal, and M. Qureshi, "Adapt: Mitigating idling errors in qubits via adaptive dynamical decoupling," in *IEEE/ACM MICRO*, ser. MICRO '21, 2021, p. 950–962.

[7] S. Niu and A. Todri-Sanial, "Pulse-level noise mitigation on quantum applications," *arXiv preprint arXiv:2204.01471*, 2022.

[8] J. Liu and H. Zhou, "Reliability modeling of nisq- era quantum computers," in *IEEE International Symposium on Workload Characterization (IISWC)*, 2020, pp. 94–105.

[9] V. Saravanan and S. M. Saeed, "Test data-driven machine learning models for reliable quantum circuit output," in *IEEE European Test Symposium (ETS)*, 2021, pp. 1–6.

[10] V. Saravanan and S. M. Saeed, "Data-driven reliability models of quantum circuit: From traditional ml to graph neural network," *IEEE TCAD*, 2022.

[11] H. Wang, P. Liu, J. Cheng, Z. Liang, J. Gu, Z. Li, Y. Ding, W. Jiang, Y. Shi, X. Qian, D. Z. Pan, F. T. Chong, and S. Han, "Quest: Graph transformer for quantum circuit reliability estimation," 2022.

[12] J. Cheng, H. Wang, Z. Liang, Y. Shi, S. Han, and X. Qian, "Topgen: Topology-aware bottom-up generator for variational quantum circuits," *arXiv preprint arXiv:2210.08190*, 2022.

[13] H. Wang, Y. Ding, J. Gu, Y. Lin, D. Z. Pan, F. T. Chong, and S. Han, "Quantumnas: Noise-adaptive search for robust quantum circuits," in *IEEE HPCA*, 2022, pp. 692–708.

[14] V. Saravanan and S. M. Saeed, "Graph neural networks for idling error mitigation," in *Proceedings of the IEEE/ACM ICCAD*, 2022.

Photon-matter quantum interfaces for scalable and networked quantum computers

Erhan Saglamyurek[1,2], Wenji Wu[1], Inder Monga[1], Alp Sipahigil[1,2], Hartmut Haeffner[2],
1. Lawrence Berkeley National Laboratory 2. University of California, Berkeley

Topic: Applications, Co-design and integration

Challenge: Quantum computation promises unprecedented computational power and capability for tackling certain problems, which cannot be solved by conventional computation. The recent technical advances have resulted in the first proto-type quantum processors with up to a few hundred qubits and gate fidelities above 99% using different platforms. Despite this impressive progress, the realization of practically useful quantum computers is challenging due to limitations on the number of qubits (several thousands high-fidelity qubits are needed) concurrently featuring high-fidelity operation, and the lack of light-based interconnection to the envisioned quantum-internet nodes. Our vision is to address these two challenges based on the tools and concepts of quantum communication, exploiting photon-matter quantum interfaces for remote entanglement between matter qubits at macroscopic distances from millimeters to hundreds of kilometers.

Opportunities: Establishing remote entanglement in both "propagating" and "stationary" forms is the key for quantum networking. This task can be accomplished with photon-matter quantum interfaces, allowing the generation and transfer of entanglement in/to the forms of photon-photon, photon-matter and matter-matter for distribution, storage and on-demand use [1]. Exploiting such capabilities of remote entanglement, photon-matter interfaces offer important opportunities for quantum computing technologies, which currently rely on advancing the quality, speed and number of local entanglement operations.

One of these opportunities is to develop modular quantum processors with photonic interconnects to overcome the scalability problem that currently all quantum computing developments face [2]. In this context, regardless of the nature and type of the platform, increasing the number of qubits in a single processing device causes more errors due to various types of cross-talks and noise. For example, for trapped ion based processors, adding more ion-qubit in the system increases the chance of cross-talks between the collective vibrational modes of ions (which are in the essence of local entangling operations), and it also reduces the speed of the gate operations, leading to more decoherence. Although neutral atom processors, relying on Rydberg interactions for local entanglement, have an advantage of putting many qubits together without such effects, trapping and addressing qubits in large arrays with high-intensity laser beams bring unavoidable cross-talks and limitations in terms of power dissipation. Similarly, as the number of qubits is increased in superconducting quantum circuits, unwanted (capacitive or microwave) coupling between qubits, higher chance of faulty or bad-quality qubits (lowering the average performance severely), and complexity of wiring the circuit to control electronics outside the cryogenic environment become major issues, restricting the qubit number to a few hundreds in the most advanced processors.

Fortunately, these limitations can be circumvented by adapting the modularity concept together with the versatility of photonic interconnect between modules. This approach leads to retain reasonably small numbers of qubits in any quantum processor module, while allowing for scale-up by linking processors through photon-mediated entanglement distribution, as shown in Figure 1. In this way, building large-scale quantum computers relies on managing photon-based quantum links between high-performance small-scale processor modules, which are within the reach of current technologies.

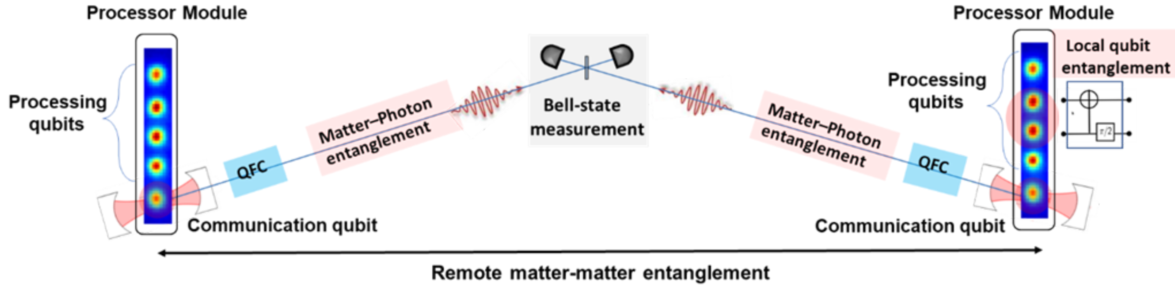

Figure 1. **Modular quantum processors with photonic interconnect**: Two processor units, each consisting of 5 qubits, are entangled via photonic Bell-State measurement, following a cavity-enhanced photon-emission of the communication qubit in each unit. The established remote entanglement may be combined with local entanglement operations for scaling up the processing capability of the units. Quantum frequency conversion (QFC) from infrared-to-telecom wavelength allows for long-distance remote entanglement over fiber-optic links.

Furthermore, photon-mediated entanglement between qubits within the same processor module can be used to enhance the qubit connectivity beyond the range of local entangling interactions, and thus such a network can reduce the number of gate operations for executing quantum algorithms [3,4]. Similarly, photon-mediated exchange of entanglement from a processing qubit to a long-lived memory qubit can be leveraged for reducing the number of physical qubits [5]. This approach becomes particularly attractive with the use of temporally and spatially multiplexed quantum memories, which can lead to novel architectures requiring substantially less number of physical qubits compared to the traditional ones [6].

While the incorporation of photon-matter quantum interfaces offers a path towards achieving scalability and reduction in the required number of physical qubits and gate operations as described above, it provides a natural way to construct a network of quantum computers [7]. In this context, quantum repeaters, which also need photon-matter interfaces, can be directly utilized for long-distant applications of quantum computers for a future quantum internet [8]. Finally, the realizations of heterogeneous architectures of different kinds of quantum computing platforms (e.g, trapped ion and neutral atom systems) would be possible via photon-assisted remote entanglement between the platforms of interest.

Assessment: The practical use of photon-matter quantum interfaces in quantum computation relies on remote entanglement to be generated at high-rates and high fidelities while storing it sufficiently long [2,9]. While the entanglement generation rates must be much faster than decoherence rates of the processing units, the fidelities must be, at least, in the range of 80% for a subsequent entanglement distillation process with optimal resources [15,16]. Long-distance applications (over kilometers) impose additional constraints to these parameters [10]; memory lifetimes must be much longer than the photon propagation time across communication links, and the entanglement generation rates must also dominate over photon loss-rates and single-mode communication rates with the supplement of temporal and/or spatial multiplexing.

The state-of-the-art experimental demonstrations have already shown the feasibility of reaching these benchmark parameters with the use of quantum processing modules [10-14]. While implementations with trapped-ion and neutral-atom based quantum processors have an inherent advantage of optical photon-interface with the possibility of telecom operation for large scale networking [12,13] superconducting qubit platforms can utilize short-range transmission of microwave photons inside their ultracold habitat (dilution fridges at ~10 mK) for building modular quantum processors [14].

Timeliness or maturity: Moving forward from proof-of-principle experimental demonstrations to actual realizations requires combining quantum communication and computing technologies, which have been seemingly two independent branches so far. In this regard, the utilization of well-established integrated photonics is critical for developing high-performance photon-matter interfaces for quantum computing. In particular, incorporation of chip-scale and robust photon cavities in the quantum processing units will be an enabling technology. In parallel, pursuing extensive theoretical research activities that explore the advantages of photon-mediated entanglement for scalability is timely important towards developing new architectures for the next generation distributed quantum computers. At the current stage, the active players of the quantum computing industry have not strategically prioritized the quantum networking aspect, but we believe that shifting some effort to this direction will accelerate the development of scalable and networked quantum computers.

References:

1. H. J. Kimble. *Nature* **453,** 1023-1030 (2008).
2. C. Monroe et al, *Phys. Rev. A* **89**, 022317 (2014).
3. A. Periwal et al, *Nature*, **600**, 630–635 (2021).
4. B. Mariennelli et al, arXiv:2303.03507 (2023)
6. É. Gouzien et al, *Phys. Rev. Lett*. **127**, 140503 (2021).
7. J. P. Covey et al, preprint at *arXiv:2304.02088* (2023).
8. H.-J. Briegel et al, *Phys. Rev. Lett*. **81**, 5932 (1998).
9. D. Hucul et al, *Nat. Phys*. **11**, 7–42 (2015).
10. V. Krutyanskiy et al, *npj Quantum Inf* **5**, 72 (2019).
11. L.J. Stephanson et al, *Phys. Rev. Lett.* **124**, 110501(2020).
12. M. Pompili et al, *Science*, **372,** 6539 (2021).
13. T. van Leet et al, *Nature* **607**, 69–73 (2022).
14. J. Qiu et al, preprint at *arXiv:2302.08756* (2022).
15. R. Rigmatulling et al *New J. Phys.* **18** 103028 (2016).
16. H. Haeffner et al, *Physics Reports*, **469,** 4 (2008).

Spectral engineering of quantum light for interconnecting different types of qubit platforms

Erhan Saglamyurek[1,2], Yiyang Zhi [2], Wenji Wu[1], Inder Monga[1], Hartmut Haeffner[2], Alp Sipahigil[1,2]

1. Lawrence Berkeley National Laboratory 2. University of California, Berkeley

Topic: Applications, Co-design and integration

Challenge: Quantum information science promises new technologies based on various quantum systems, each offering an unique advantage. Combining these technologies, a future quantum internet will likely rely on distant nodes, composed of different types of quantum matter platforms, interfaced with optical photons [1]. Optical interconnection of such nodes over long distances is challenging due to the incompatibility of the platforms' wavelengths with low-loss spectral window of fiber-optic and free-space channels as well as mismatch between the wavelength and bandwidth of the quantum systems [2]. While the traditional tools of quantum photonics offer efficient optical wavelength-conversion techniques over a wide spectral range for alleviating the wavelength mismatch and lossy transmission, there is a vast temporal wavepacket-length mismatch between the commonly studied qubit platforms. Therefore, it is necessary to find solutions to match the length of single-photon wavepackets via coherent conversion processes which optimize the interference visibility in heterogeneous quantum networks [3].

Opportunities: Well-known qubit platforms, including trapped ions, neutral atoms, solid-state defect centers, feature optical photon interfaces in the visible and infrared range with Fourier-limited bandwidths from several kilohertz to terahertz (microseconds to picoseconds in the time domain). Such a spectral mismatch between the systems may be circumvented by coherently stretching or compressing temporal profiles of photons while equalizing their carrier frequencies in the same process or with an independent process following/preceding the bandwidth conversion.

To this end, two different approaches have been proposed and experimentally demonstrated for optical pulse compression/stretching. The first one relies on the concept of a time lens, requiring a highly dispersive optical medium in conjunction with time-dependent phase modulation, see Fig. 1a [4-8]. The role of the dispersive medium is to introduce a frequency dependent group-delay (resulting in stretching of pulse duration/compressing of bandwidth) while the phase modulation allows for compensating or re-distributing the spectral/phase profile. This approach has been implemented using various dispersive platforms (e.g., engineered optical fiber, chirped fiber grating) together with phase modulation methods relying on electro-optic elements and three/four wave mixing processes in non-linear media. The minimum compressed bandwidth (or the maximum achievable compression/stretching factor) is given by the amount of the delay-time difference in the dispersive media, which can, in practice, lead to conversion between the GHz and THz bandwidths without too much loss. This feature brings the possibility of reducing the bandwidth of the inherently wide spectrum of photons from certain solid-state-based platforms to a level closer to the bandwidth range of atomic platforms (Fig. 1c).

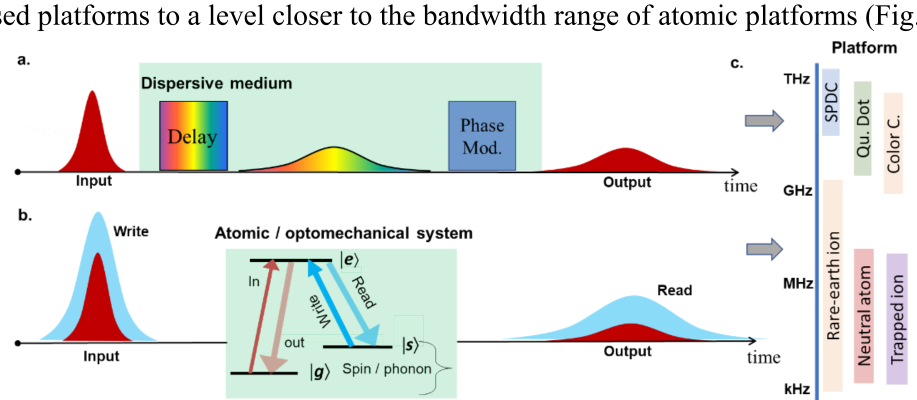

Figure 1. **Conceptual descriptions of bandwidth conversion**: **a.** An input signal is sent to a dispersive medium, which introduces a frequency dependent delay and hence stretches the duration of the pulse. Phase modulation recovers the original frequency/phase profile, resulting in the bandwidth compression. **b.** A three-level system, comprised of an optical level |e> and two ground spin/ vibrational levels |g> and |s>, is used to map coherence from an input signal to spin/phonon excitations via an optical field (write) .After a certain time, the stored state is converted back to an optical signal via another optical field (read) with a duration longer than the input, resulting in stretched temporal profile. **c.** Typical operation bandwidth ranges of widely employed matter and photonic qubit platforms. SPDC : Spontaneous parametric downconversion.

The second approach uses atomic or optomechanics platforms via the implementation of a light-matter interaction protocol (e.g. electromagnetically\optomechanically induced transparency) for reversible mapping of optical coherence to spin or phononic (mechanical) modes, respectively, as shown in Fig 1b [9-12]. In this approach, an input optical signal with a certain initial bandwidth is coherently stored in the form of spin/phonon excitations. After a desired time, the stored coherence is mapped back to an output optical signal with a target bandwidth that can be set to be wider or narrower than that of the input. An inherent advantage of this approach is that the very same mechanism can also serve as a quantum memory as well as a frequency conversion unit. Furthermore, dynamic controlling the bandwidth allows for compressing and stretching in a flexible manner from the kHz to GHz regime [8,12] provided that the light-matter coupling is sufficiently large, both in the adiabatic as well as in non-adiabatic interaction regime. Optomechanical systems provide additional flexibility for tuning optical wavelengths to arbitrary values and thus renders the wavelength matching with the widely studied qubit platforms.

Assessment: Efficiency, fidelity, and initial/target bandwidths are the main parameters that determine the performance of the bandwidth conversion process. While the dispersion engineering approaches suffer from propagation losses for long group-delays required for GHz-range target-bandwidths, near unity efficiency is still within the reach for larger bandwidths. The efficiency for atomic ensemble- and optomechanical-based approaches is mainly determined by light-matter coupling strengths and can approach unity with large atomic densities and fabrication of nano-optomechanical devices, respectively.

Fidelity is another critical parameter that may be limited by photonic noise arising from strong optical fields to mediate the conversion process. In this respect, some of the dispersion-based approaches do not involve a pump field in the process and hence promise high-fidelity operation. In addition, optomechanical devices need to be cooled down to their ground states in order to eliminate thermal noise of mechanical modes, which may be technically challenging due to heating effects arising from strong pump fields. While ensembles of cold atoms are considered to be free from thermal noise, suppressing the pump-field noise may be difficult due to unavoidable four-wave mixing effects at large initial or target bandwidths (~GHz).

Finally, the initial and target bandwidths should be as far separated, especially for interfacing solid-state and atomic qubit platforms, which typically operate at several GHz and kHz bandwidths, respectively. Fulfilling this demand requires a conversion system, concurrently featuring a large time-bandwidth product (100>) and a large dynamic-operation range (e.g., MHz-GHz). The challenge is to combine such a large time-bandwidth product and dynamic range with an efficient (>90%) and high-fidelity operation (>90%).

Timeliness or maturity: The state-of-the art experiments have already reached the above-described benchmarks in separate implementations. However, significant work remains to be done for practical converters that are able to interface qubit platforms without degrading their performance. Developing such technology is both important and timely as several national and world-wide initiatives have started establishing quantum network test-beds based on different qubit technologies. In this direction, following a systematic evaluation of the existing techniques, a multidisciplinary effort needs to be put into engineering customized systems that can interface specific platforms (e.g., trapped ions and color-centers). This effort could be further extended towards the development of more general quantum optical processors with multi-tasking capability, combining frequency-bandwidth conversion, storage and multiplexing.

References:

1. N. Gisin et al. *Nat. Photon* **1,** 165-171 (2007).
2. C. Simon , *Nat Phot.* **11**, 678-680 (2017).
3. F. Sosnicki et al, arXiv:2211.05112v1 (2022).
4. J. Lavoie et al, *Nat. Photon.* **7,** 363–366 (2013).
5. M. Karpinski et al *Nat Photon*. **11**, 53 (2017).
6. I. Agha et al, *Opt.. Lett*. **33**, 1047 (2008).
7. C. Josh et al, *Optica*, **9**, 364 (2022).
8. E. Saglamyurek et al, *New J. Phys*. **16** 065019 (2014).
9. M. Hosseini et al, *Nature* **461**, 241 (2009).
10. J. L Le-Gout et al *Phys. Rev. A*. **80**, 012320 (2009).
11. J. T. Hill et al, *Nat. Comm*, **3,** 1196 (2012).
12. E. Saglamyurek et al, *Nat Photon,* **12** 774 (2018).

Visualizations for Uncertainty Analysis of Quantum Application Outputs

PRIYABRATA SENAPATI, Kent State University, USA
TUSHAR M. ATHAWALE, Oak Ridge National Laboratory, USA
DAVID PUGMIRE, Oak Ridge National Laboratory, USA
QIANG GUAN, Kent State University, USA

TOPICS

Visualization of uncertainty for noisy states in quantum computing, Quantum error mitigation, Impact of noise, Quantum computing application and performance modeling

CHALLENGE

Noise in the quantum computers today presents a main challenge to the users despite the rapid progress. We have seen two techniques, quantum error correction and quantum error mitigation, for addressing the noise so that the quantum computers provide stable output for high performance computing. Quantum error correction requires a large number of qubits [Acharya et al. 2023] which is not present in quantum computers today. To make it worse, errors could grow with the number of qubits. The other technique, quantum error mitigation, approaches the noise problem in such a way that the effect of the noise in the output is mitigated. In order to make our quantum computers more viable, we need to mitigate the noise in such a way that the impact of the noise in the output is less critical. Recent work [Ravi et al. 2023] has demonstrated reducing the impact of error by separating negative gradients and positive gradients configurations, but we have not come across any works that has attempted to isolate noisy and non-noisy qubits. One of the challenges in quantum computing is reproducibility. Unlike reproducibility in classical computing, quantum computing is not reproducible due to the dynamic changes in noise landscape. Quantum machine learning algorithms when tested on a quantum computer yields varying outputs across various quantum computers.

OPPORTUNITY

Currently, there is no adequate framework to convey noisy and non-noisy basis states in quantum computers. Such a lack of framework can be addressed through researching novel visualization and analysis techniques that can effectively and efficiently communicate the noisy and non-noisy states to the users. One of the reasons for noise in quantum computers is noisy qubits. Not all physical qubits are noisy, and noisy qubits could be differentiated from non-noisy qubits. Visualization techniques can help in simplifying the task of decoding noise in quantum computers, and without access to visualization tools, it will be difficult to understand noise in quantum computers. As illustrated in Figure 1a, our prototype visualization can aid in analysis of noise in 128 basis states for 7 qubits quantum computer IBM Nairobi. Further, noise distribution across different states can be efficiently compared through visualization of pairwise KL-distance among various quantum states, as shown in Figure 1b. We hypothesize that visualizing the noise in basis states is a critical step in the direction for a meaningful quantum error mitigation. To achieve this goal, novel statistical and classical machine learning techniques are required to be researched for visualization and analysis

Authors' addresses: Priyabrata Senapati, Kent State University, 1300, Lester A Lefton Esplanade, Kent, Ohio, USA, psenapat @kent.edu; Tushar M. Athawale, Oak Ridge National Laboratory, Tennessee, USA; David Pugmire, Oak Ridge National Laboratory, Tennessee, USA; Qiang Guan, Kent State University, 1300, Lester A Lefton Esplanade, Kent, Ohio, USA.

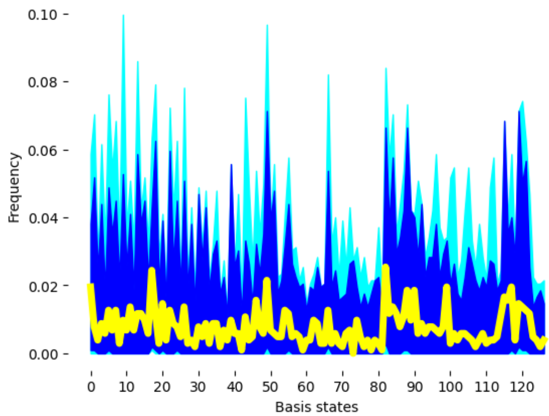

(a) Executed QML jobs from IBM Nairobi

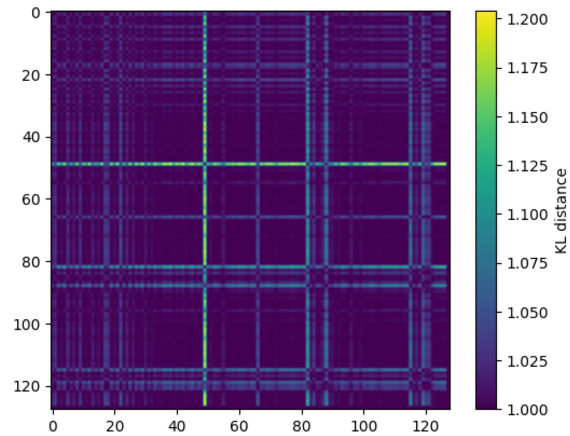

(b) KL-distance for executed QML jobs from IBM Nairobi

Fig. 1. Importance of visualization for analysis of noise in quantum computers. (a) Statistical visualization of state frequencies provides an insight into the median frequency in yellow. The blue bands indicate uncertainty around the median observations. (b) The heatmap of a pairwise KL-distance between state distributions highlights (yellow) states with significantly different distributions compared to others.

of high-dimensional quantum state data. Further, the complexity of analysis of quantum states grows exponentially with the number of qubits. Thus, researching novel scalable visualizations is essential to conveying noise in quantum computers for large number of qubits.

ASSESSMENT

Once noisy qubits are identified using visualizations, it is important to test the quantum machine learning model on non-noisy qubits and then use a confusion matrix to measure the accuracy of the model. A confusion matrix summarizes the performance of a classification algorithm. The confusion matrix of quantum computer data should be compared against the confusion matrix obtained from ground truth quantum-simulator data. This way a sound assessment can be carried out to ascertain our error mitigation methodology.

TIMELINESS

We have seen a rapid development [IBM 2020] in the past few years building quantum computers. It was only a few years ago when IBM allowed the general public to use its quantum computers with fewer than 7 qubits. Today, we see companies like IBM have started offering people access to 400 plus qubits quantum computer. It is a high time that we work around with the noise in quantum computers to make it a highly stable and viable tool for high performance computing. With our idea of visualizing noise in quantum computers, users could potentially avoid using certain noisy qubits that may effect the output.

REFERENCES

R Acharya et al.. 2023. “Suppressing quantum errors by scaling a surface code logical qubit.” *Nature*, 614, 7949, 676–681.

IBM. 2020. “IBM’s Roadmap For Scaling Quantum Technology.” *https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap/*. Accessed: 2020-09-30.

Gokul Subramanian Ravi, Kaitlin Smith, Jonathan M Baker, Tejas Kannan, Nathan Earnest, Ali Javadi-Abhari, Henry Hoffmann, and Frederic T Chong. 2023. “Navigating the dynamic noise landscape of variational quantum algorithms with QISMET.” In: *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, 515–529.

Distributed Quantum Processing via Integrated Regional Quantum Networks Using Free-Space Optical Links

Harry Shaw (NASA, harry.c.shaw@nasa.gov), LaVida Cooper (NASA, lavida.d.cooper@nasa.gov), Gabriella Carini (BNL, carini@bnl.gov), Dimitrios Katramatos (BNL, dkat@bnl.gov), Kathy-Anne Soderberg (AFRL, kathy-anne.soderberg@us.af.mil), Eden Figueroa (Stony Brook U./BNL, eden.figueroa@stonybrook.edu)

Topic: Quantum Networking Models and Applications

1. Challenge: Integration of regional quantum networking infrastructure.

The landscape for scalable quantum processing is rapidly evolving with the development of quantum computers and the emergence of future science enablers, such as distributed quantum computing and hybrid computing scenarios where quantum and classical computing assets work cooperatively. Quantum entanglement distribution is also advancing with multiple regional quantum networks and testbeds maturing across the nation. These quantum networking testbeds typically include multiple access points with a heterogeneous mixture of capabilities interconnected by fiber spanning anywhere from tens to hundreds of kilometers; many such networks incorporate a classical networking framework. Alongside these quantum connectivity advancements in the infrastructure, the maturation of quantum sensing technology presents an opportunity to expand the body of quantum networking use-cases that are presently possible. Along these lines, bridging the present collective progress in quantum computing and quantum networking will enable us to develop new use-cases. One such use-case is enabling future quantum computing applications via integrated regional quantum networking infrastructure in neighboring locations like New York and Maryland, both states are home to robust classical and quantum networking capabilities. To enable near-term distributed quantum computing applications via the practical integration of regional quantum networking infrastructure, free-space optical links via satellite communications will be essential to avoid the exponential loss commonly observed with fiber implementations [1]. There is both an urgent need and a practical justification for pursuing such space-based integration scope now.

2. Opportunity: Quantum Applications using space-based quantum networking.

It is paramount to develop applications that can take advantage of the special nature and non-locality of entanglement. Present applications include blind quantum computing, unprecedented space-science data processing algorithms to advance fundamental physics research, and the ability to break all current forms of cryptography that are not post-quantum hardened. Additionally, applications in quantum sensing are already emerging. Ultimately, the advantages that exist for classical space communications will be enhanced by the novel resources and insights that these quantum applications will provide. To effectively enable such quantum applications over long distances, legacy classical optical communications must be adapted for space-based quantum networking scenarios. In order to achieve first realizations of space-based quantum links, it will be necessary to develop new quantum-ready components, to design integrated classical/quantum systems performance, to envision novel concepts of operation and enhanced end-to-end data flow paradigms, and to address flight/orbital dynamics challenges and precision pointing, acquisition and tracking needs for ground segment telescopes. In the near-term we can envision an extension of classical networks to include the ability to pass quantum entangled data through the network. In this scenario, the “heavy lifting” in terms of data rate, and performance should be confined to the classical channels. The quantum channel should be used for functions such as entanglement and entanglement swapping. Additionally, augmenting the Open Systems Interconnection (OSI) protocol stack will be required to accommodate the inclusion of quantum data flow in the network. Furthermore, network security, including quantum and post-quantum cryptography methods must be designed into all the network elements. Further progress can be achieved by further developing quantum devices such as quantum memories and entanglement sources, which exist primarily in laboratories and testbeds, into miniaturized systems that can be implemented in spacecraft with a reasonable size, weight, power, and cost (SWAP-C) within a reasonable timeframe. These challenges can be investigated at the laboratory level now and national programs should incentivize the near-term demonstration of these solutions in the context of space-based projects. The future of such space-based quantum-network missions will also depend on the ability of legacy optical communications commercial off the shelf (COTS) sources, detectors, and optics to be used with little or no modifications. Therefore, an intentional commercialization strategy must be an additional spin-off of lab- scale and/or larger-scale quantum-networking research and national investment programs. By reducing and eventually removing the overhead of non-recurring engineering for key quantum networking components and systems designed to perform over target functional ranges, quantum local area networks around the country will be able to scale up their network capacity and availability in an efficient manner and thus accelerate their readiness for integration and intracontinental operations aided by space-based optical communications links. The ability to space qualify key components and parts is also a challenge to overcome.

3. Assessment: Establishing LEO or MEO space- based optical quantum communication.

A targeted and nationally supported program is needed to enable the practical tasks associated with defining, developing, testing, evaluating, and operating a free-space connection to interface regional quantum networking access points where quantum computing assets can either be remotely leveraged or physically co-located to implement science enabling distributed quantum processing tasks. Success of such a program would be defined as the ability of the New York and Maryland collaborations to communicate via network protocols that include exchange of quantum entanglement via a LEO or MEO space- based optical communications flight asset with ground terminals coupled to terrestrial quantum networking infrastructure in each state [2].

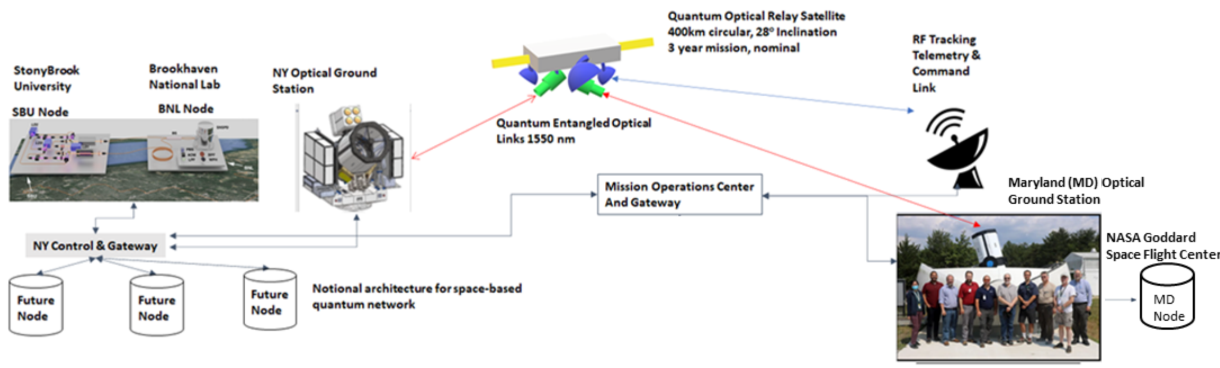

4 Maturity: Achieving *operational* competency in spaced-based quantum networking.

It took several decades to mature the current world-wide internet from its origins in the ARPANET program. This maturation process was accelerated by the development of Ethernet, TCP/IP, UDP, and other data communication protocols. Simultaneously, the physical layer infrastructure had to be laid into place. The same situation exists today regarding the development of spaced-based quantum networking and distributed quantum computing. The United States has spent years maturing key quantum components, devices, systems, executing studies, analysis, while also building quantum-enabling infrastructure. Furthermore, there is a wealth of classical networking infrastructure and approaches widely used now that could serve as a guide to where adaptations are needed. The ability to form spatially diverse worldwide quantum optical communications networks will involve free space communications through the atmosphere to earth orbiting relay satellites, which can use our well established abilities to build and operate earth orbiting relay satellites and ground stations and to operate effective mission control. Within this framework, collaborations between domestic organizations with government support should help achieve national *operational* competency in spaced-based quantum networking. The impact of success would be a tangible demonstration of intracontinental quantum computing along the east coast of the United States, realizing a subset of the objectives set forth in the national quantum strategy.

[1] Science 356, 1140–1144 (2017), [2] EPJ Quantum Technology 5, 6 (2018).

Focus beyond Surface Code for Error-Corrected Quantum Advantage

Ruslan Shaydulin*§ Michael A. Perlin† Ciarán Ryan-Anderson‡
Shouvanik Chakrabarti* Marco Pistoia*
*Global Technology Applied Research, JPMorgan Chase, New York, NY, USA
†Inflegtion, Chicago, IL, USA
‡ Quantinuum, Broomfield, CO, USA

Topic Compilation, error correction and mitigation.

Challenge Quantum error correction (QEC) is known to introduce significant overhead. A deep understanding of this overhead is essential for identifying quantum algorithms that can provide computational speedups over classical state-of-the-art. For example, a series of recent results suggest that the overhead of surface code makes quadratic speedups unlikely to provide computational advantage for early fault-tolerant superconducting quantum computers [1, 2]. However, the optimization of QEC protocols for non-fixed-qubit architectures is a far less explored subject and, thus, these results may not be applicable for such architectures. Hardware platforms such as trapped ions and cold atoms have hardware features, most notably all-to-all or otherwise beyond 2D nearest-neighbor connectivity, that unlock alternative avenues for achieving fault tolerance [3, 4]. These capabilities introduce two challenges: (1) quantifying the extent to which such hardware features reduce the overheads of QEC, and (2) understanding the extent to which these reduced overheads enable quantum advantage with quadratic or otherwise low-degree polynomial algorithmic speedups. A precise understanding of what is and is not a "promising" speedup is of central importance to the quantum computing community, as it informs the directions that researchers pursue.

Opportunity Recent theoretical developments in quantum low-density parity-check codes (LDPC) codes (e.g., [5–7]) as well as experimental advancements in non-fixed-qubit hardware (e.g., [8, 9]) create an opportunity to re-evaluate common assumptions about the overheads of fault tolerance and, in turn, the induced limitations for achieving practical algorithmic speedups with early fault-tolerant devices. Codes leveraging long-range interactions have been shown to reduce qubit (space) overheads by orders of magnitude compared to the surface code [4] and can similarly reduce time-complexity overheads [6]. For example, a 3D qubit topology enables the execution of codes with "single-shot" fault tolerance [10], meaning that the time overhead of QEC does not grow with the code size, in contrast to the typical implementation of the 2D surface code. Such features of quantum error-correcting codes have been traditionally sidelined or otherwise considered merely speculative, reflecting a perspective largely centered on fixed-qubit architectures with nearest-neighbor planar connectivity. In addition, Ref. [2] found that the classical processing overhead of decoding is a limiting factor for quantum speedups based on surface-code QEC. However, the longer coherence times of trapped ion and neutral atom platforms relax the requirements for classical decoders, as evidenced by the experimental realization of real-time error correction with trapped ions [11, 12]. In general, longer coherence times relax the latency constraints between interacting classical and quantum classical devices facilitating conditional branching and adaptive quantum circuits allowing for highly dynamical QEC protocols [13, 14] as well as allowing for the reduction of needed syndrome extraction rounds during logical memory [15], all of which further reduce the overhead of QEC. Combined, advances in theory and hardware present an opportunity to re-evaluate the common wisdom about the overheads of QEC and overcome the bias towards fixed-qubit architectures present in the resource estimation literature.

§ Corresponding author. Email: ruslan.shaydulin@jpmchase.com

Assessment A successful program would provide a quantitative, full-stack evaluation of the QEC overhead of novel LDPC codes enabled by different hardware approaches, as well as the algorithmic speedups compatible with such overheads. This evaluation should include a detailed, hardware-specific resource analysis for implementing novel LDPC codes, taking into consideration such factors as, e.g., the overheads of moving atoms with optical tweezers or shuttling ions in a QCCD architecture. Other factors such as parallelization, physics-aware noise models, and concrete decoding schemes should be taken into account since the performance of codes and classical compute time overhead will be affected and, thus, influence the choice of QEC codes and qubit-time overheads. To identify and mitigate risks, the resource analysis should evaluate the impact of various assumptions about hardware improvements and future capabilities such as multi-qubit gates, degree of connectivity needed, and shuttling times. To translate the resource analysis into practical implications for quantum algorithms, successful programs should additionally identify the conditions that QEC codes have to satisfy to make certain kinds of speedups (e.g., quadratic) practical, which includes identifying no-go results for algorithmic speedups from fundamental limitations of QEC.

Timeliness or maturity The maturity of surface code and vendor roadmaps for scaling fixed-qubit architectures have enabled end-to-end resource analyses, connecting the overheads across the quantum stack to produce high-confidence conclusions about the algorithmic speedups necessary for a quantum computational advantage. In a similar vein, LDPC codes have recently seen major theoretical breakthroughs, including the discovery of so-called "good" LDPC codes with constant rate (logical per physical qubits) and linear distance (growing with the number of physical qubits) [5], and the revival of LDPC code concatenation as a viable pathway to scalable QEC [6, 7]. These developments have the potential to dramatically reduce the physical overhead required for fault-tolerant QEC. At the same time, trapped ion and cold atom devices with 10s–100s of qubits are coming online, and vendors are laying out concrete and realistic roadmaps for future hardware developments. These devices have demonstrated critical features such as low memory errors, high-fidelity mid-circuit measurements and qubit resets, low crosstalk, and dynamic circuits [9]. The developments in non-fixed-qubit hardware platforms, and the flexibility that these platforms provide for implementing a broad range of quantum error-correcting codes, makes end-to-end resource analyses of viable quantum algorithmic speedups not only possible but *urgently necessary* to guide further hardware development efforts and inform near-term demonstrations of fault-tolerant QEC primitives. Further, such analysis is crucial to the field of quantum computation to inform quantum algorithm developers whether hardware with less stringent connectivity constraints provides a route for quadratic or low-degree polynomial speedups for future large-scale fault-tolerant computation, or whether it is better to shift research to other algorithms.

[1] R. Babbush et al., PRX Quantum **2**, 010103 (2021).
[2] E. Campbell et al., Quantum **3**, 167 (2019).
[3] L. Z. Cohen et al., Science Advances **8**, eabn1717 (2022).
[4] D. Litinski et al., (2022), arXiv:2211.15465.
[5] P. Panteleev et al., in Proc. of the 54th Symposium on Theory of Computing, STOC 2022 (2022).
[6] H. Yamasaki et al., (2022), arXiv:2207.08826.
[7] C. A. Pattison et al., 2023, arXiv:2303.04798.
[8] D. Bluvstein et al., Nature **604**, 451–456 (2022).
[9] J. M. Pino et al., Nature **592**, 209–213 (2021).
[10] H. Bombín, Physical Review X **5**, 031043 (2015).
[11] C. Ryan-Anderson et al., Physical Review X **11**, 041058 (2021).
[12] C. Ryan-Anderson et al., 2022, arXiv:2208.01863.
[13] C. J. Trout et al., New Journal of Physics **20**, 043038 (2018).
[14] T. Tansuwannont et al., Phys. Rev. A **101**, 012342 (2020).
[15] X. Xu et al., New Journal of Physics **20**, 023009 (2018).

Disclaimer

Advanced noise handling capabilities: from error mitigation to error correction

Andrew Cross, Sarah Sheldon, and Ewout van den Berg
IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA
(Dated: May 1, 2023)

Topics: Error Mitigation and Error Correction, codesign, applications

A. Challenge

Addressing noise in quantum systems is critical for realizing useful quantum computing. The development of quantum error correction (QEC) and discovery of codes with high thresholds has given a plausible trajectory toward fault tolerant quantum computing. Well-studied codes like the surface code provide a guide for hardware development, setting specifications for connectivity and error rates. Despite this, fault tolerance is out of reach of near term quantum computers. At the same time, quantum systems are now available at a size that is beyond exact classical simulation. These devices motivate research on near term quantum algorithms and the search for problems that could be solved on such devices.

Current research typically takes one of two tracts: demonstration of error correction subroutines such as preparation of logical states, or near term applications research that disregards noise in quantum circuits. But given that we have devices with over 100 qubits and capable of circuits depths that go beyond classical diagonalization, we need to find ways to extract useful and reliable measurements from these noisy devices in reasonable runtimes [1].

In the near term this requires error mitigation techniques and problems that involve expectation value measurements. Problem-agnostic error mitigation techniques such as Probabilistic Error Cancellation (PEC) and Zero Noise Extrapolation (ZNE) [2, 3] have now been demonstrated for simple test circuits [4, 5]. Over a longer time horizon, we need more practical ways of implementing error correction; we need better error correction codes with lower overheads.

There are several challenges associated with the continuous improvement of error handling techniques and identifying useful problems capable of being solved on near term quantum computers:

- *Verifying the outputs of error mitigated circuits*: when circuit sizes become large enough that we cannot simulate outputs classically, we need alternative methods to give us confidences that mitigated results are correct.
- *Constructing accurate noise models*: error mitigation methods such as Probabilistic Error Cancellation (PEC) and Zero Noise Extrapolation (ZNE) rely on having good models of the device noise and making reasonable assumptions about the nature of the noise. As device errors become smaller, validating these assumptions is critical.
- *Reducing overheads*: PEC offers guarantees that the mitigated estimator is unbiased as long as the noise model is accurate, but it comes at an exponential sampling overhead cost. ZNE appears heuristically to have lower overheads but at the risk of introducing larger bias. Extensions or simplifications of these techniques could have significant impacts on overall runtime of error mitigated quantum circuits. In parallel, we need further theory exploration on QEC codes to bring down the number of physical qubit required.
- *Finding use cases*: for these techniques to be useful in the near term, we need to find problems that are amenable to the types of circuit structures and observables that can be mitigated. The availability of quantum hardware and capabilities that can run mitigated circuits should lead to new research questions in quantum computational science.

B. Opportunity

The possibilities for novel research abound in this new frontier for quantum computing. Continuing to improve the reliability of mitigated measurements through more accurate noise modeling and error mitigation techniques with lower sampling overheads remains the task of quantum information researchers, but there are opportunities for interdisciplinary collaborations. Accessing larger circuit sizes will likely spur on classical algorithms researchers to compete and drive development of classical approximation methods like tensor network simulations. Dedicated effort in this area would benefit quantum research by both allowing us to verify quantum results in regimes that are classically simulable and to identify circuit conditions that are challenging classically and therefore useful to study with quantum computers.

The intersection of quantum computational science and applications research with quantum capabilities development also provides opportunities for collaboration. Connections between these researchers should lead to feedback loops where better quantum hardware and error mitigation techniques inspires new use cases, which in turn help define classes of problems that should be targeted for near term value.

Optimizing these error handling tools and making them accessible to users requires software efforts targeting multiple levels of the stack. Software needs include efficient compilation of circuit samples, integrated noise resilience tools, accessible diagnostics, and automation for application-focused end users.

We now have a tradespace to explore within error mitigation and error correction techniques – we can look at how best to combine the outputs of circuits in classical postprocessing, or introduce encoded qubits to mitigated frameworks. This opens up a new avenue of research for quantum error correction before fault tolerance. Are there problems we could address sooner if we could lower effective errors rates through combinations of error mitigation and error correction?

Ultimately we want to chart a path towards fault tolerance, but this path can consist of gradual improvements in error handling overhead. The field is now seeing a burst of activity in new codes research, and early work on good quantum low-density parity check (LDPC) codes indicates that there are opportunities to reduce resources required for error correction [6, 7]. For the long term outlook, we should be aiming to codesign hardware with more efficient error correcting codes.

C. Assessment

We expect that the combination of quantum hardware with over 100 qubits and capabilities for extracting reliable measurements will naturally open up these new areas of research. We should see applications development focus on problems that can be addressed at this scale. Full stack quantum systems that contain the tools for providing noise mitigated estimators of observables should stimulate the search for use cases.

Success in this effort would include extensions of current error mitigation methods, a growing collection of example use cases that are at or beyond the threshold for quantum advantage, and an outline of the error rates and capabilities required to access more complex problems in the future. We should see advances on existing error mitigation techniques become integrated into software stacks and become more seamlessly integrated into workflows.

D. Timeline and maturity

PEC and ZNE have already been tested experimentally and shown to give reliable results and dramatic improvements over unmitigated data. While these methods do not apply to the many quantum algorithms that require single shot data, these results show that one can obtain good values from noisy quantum computers when limited to expectation values. These types of measurements correspond to some of the most attractive early use cases of quantum computers like measuring properties of quantum systems for chemistry or materials science.

As the performance and scale of commercially available devices keeps improving, we should be investigating what we can do with these systems now and how we can extract the highest quality data. We are beginning to see extensions and combination of existing error mitigation techniques in practice [8, 9]. We have theory proposals for frameworks that use error mitigation within an error corrected framework [10] and proof of concept demonstrations of techniques that borrow subroutine from error correction to go beyond averaged measurements [11]. Research into noise handling capabilities before fault tolerance is already bearing fruit, and it will only become more necessary as we pursue useful applications of quantum computing in the coming years.

Authors email: ssheldo@us.ibm.com

[1] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, The future of quantum computing with superconducting qubits, Journal of Applied Physics **132**, 160902 (2022).

[2] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Physical Review Letters **119**, https://doi.org/10.1103/physrevlett.119.180509 (2017).

[3] Y. Li and S. C. Benjamin, Efficient variational quantum simulator incorporating active error minimization, Phys. Rev. X **7**, 021050 (2017).

[4] E. van den Berg, Z. K. Minev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse pauli-lindblad models on noisy quantum processors (2022), arXiv:2201.09866 [quant-ph].

[5] Y. Kim, C. J. Wood, T. J. Yoder, S. T. Merkel, J. M. Gambetta, K. Temme, and A. Kandala, Scalable error mitigation for noisy quantum circuits produces competitive expectation values, Nature Physics 10.1038/s41567-022-01914-3 (2023).

[6] N. P. Breuckmann and J. N. Eberhardt, Balanced product quantum codes, IEEE Transactions on Information Theory **67**, 6653 (2021).

[7] P. Panteleev and G. Kalachev, Asymptotically good quantum and locally testable classical ldpc codes (2022), arXiv:2111.03654 [cs.IT].

[8] M. Urbanek, B. Nachman, V. R. Pascuzzi, A. He, C. W. Bauer, and W. A. de Jong, Mitigating depolarizing noise on quantum computers with noise-estimation circuits, Phys. Rev. Lett. **127**, 270502 (2021).

[9] B. McDonough, A. Mari, N. Shammah, N. T. Stemen, M. Wahl, W. J. Zeng, and P. P. Orth, Automated quantum error mitigation based on probabilistic error reduction, in *2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)* (IEEE, 2022).

[10] C. Piveteau, D. Sutter, S. Bravyi, J. M. Gambetta, and K. Temme, Error mitigation for universal gates on encoded qubits, Physical Review Letters **127**, 10.1103/physrevlett.127.200505 (2021).

[11] E. van den Berg, S. Bravyi, J. M. Gambetta, P. Jurcevic, D. Maslov, and K. Temme, Single-shot error mitigation by coherent pauli checks (2022), arXiv:2212.03937 [quant-ph].

High Fidelity Noise-Tolerant State Preparation of a Heisenberg spin-1/2 Hamiltonian for the Kagome Lattice on a 16 Qubit Quantum Computer

Wladimir Silva. Dept. of
Computer Engineering Raleigh
NC. (wsilva@ncsu.edu)

***Abstract*—This work describes a method to prepare the quantum state of the Heisenberg spin-1/2 Hamiltonian for the Kagome Lattice in an IBM 16 qubit quantum computer with a fidelity below 1% of the ground state computed via a classical Eigen-solver. Furthermore, this solution has a very high noise tolerance (or overall success rate above 98%). With industrious care taken to deal with the persistent noise inherent to current quantum computers; we show that our solution, when run, multiple times achieves a very high probability of success and high fidelity. We take this work a step further by including efficient scalability or the ability to run on any qubit size quantum computer. The platform of choice for this experiment: The IBM 16 qubit transmon processor ibmq_guadalupe using the Variational Quantum Eigensolver (VQE).**

Keywords— Quantum, VQE, Kagome

I. WHY THE KAGOME LATTICE AND QUANTUM?

Kagome lattice is an arrangement of atoms in 2-D pattern commonly found in minerals. This lattice is studied in condensed matter physics because of its geometry: When atoms in the lattice are anti-aligned by spin, the triangular shape creates a frustration (right side of figure 1), that is, the atom on the right doesn't know what to do. This type of frustrated system is believed to be of great importance in the study of spin liquids and superconducting materials at high temperatures [1]. Modeling this type of frustrated system requires a tremendous amount of memory, and that's where a quantum computer can help.

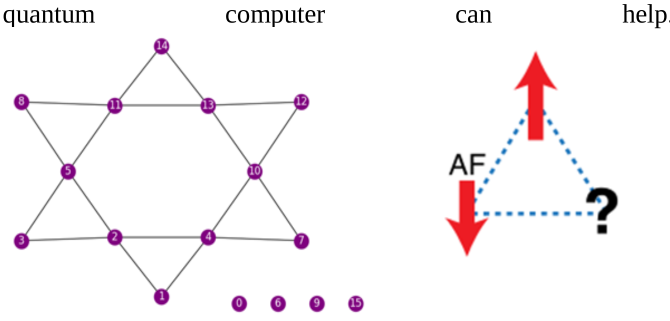

Figure 1 Kagome lattice unit cell (left). On the right, an example of a frustrated system with atoms anti aligned by spin.

II. THE MASSIVE POWER OF QUANTUM

For a long time, physicists have been modeling the ground and excited states of molecules in search for better materials and superconductors at high temperatures. A few years ago, supercomputers were the only game in town because of the gargantuan memory requirements to model even the simplest of molecules. That changed after 2018 when IBM opened up their new state of the art quantum computers to the scientific community. Only a quantum computer has the exponential power to tackle the state preparation of heavy elements: Consider a 400 qubit processor. It is capable of consuming 2400 states in parallel, far ahead of the capabilities of most supercomputers in existence. Coupling this fact with the advantages provided by quantum mechanics: superposition of states and entanglement, here we have all the tools we need to model these heavy elements.

III. STATE PREPARATION IS TRIVIAL, MANAGING NOISE IS NOT

Our experiment begins with a 4 ingredient recipe: The Variational Quantum Eigen Solver (VQE), an Ansatz (Initial state), a Hamiltonian, and the optimizer.

VQE is the de-facto algorithm of choice for modeling ground states (see figure 2). It contains three parts: An Ansatz or initial quantum circuit made of a set of rotation gates over the Y, Z axis of the Bloch Sphere parameterized by random angles. The Ansatz works with a cost function whose task is to evaluate the Hamiltonian to produce an energy value. A classical optimizer attempts to minimize this energy, by updating the angles of the Ansatz using sophisticated techniques: gradient descent, heuristics, and others. At the end, this sequence repeats for a number of cycles, ideally reaching the ground state. The energy values are collected to produce a final plot of the energy minimization process. These values are compared with a classical eigensolver to estimate the fidelity (or accuracy) of the experiment.

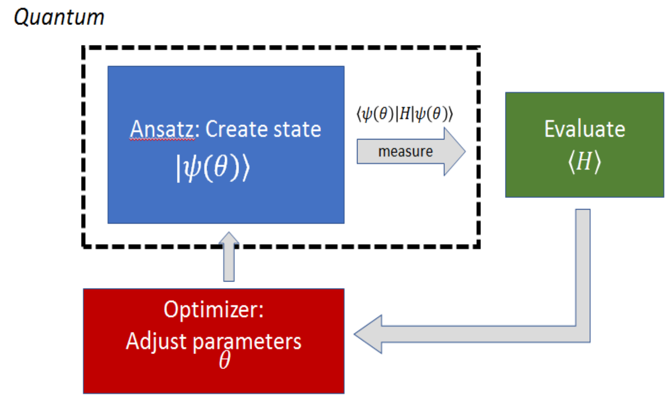

Figure 2 Layout of the VQE algorithm including the Ansatz, cost function and optimizer. Note that the Ansatz is the only piece executed in a quantum computer, the rest are classical steps.

There are two things to consider when picking an (Ansatz, Optimizer) combo which will determine the shape of the final curve:

1. **The combination of Ansatz, Optimizer:** It produces a different curve type. Consider figures 3 and 4 which show different shape types from experimental results.
2. **Number of cycles in the optimizer:** Some optimizers default to an excessive number of cycles. In figure 5, Univariate Marginal Distribution Algorithm (UMDA) defaults to 1200 cycles. This will balloon the execution time of the experiment, and worst of all, it will increase the odds of failure by memory exceptions or bugs in the server side. As a matter of fact, the experimental result from figure 5 was run in IBM's 27 qubit Cairo processor and took more than 5h to run.

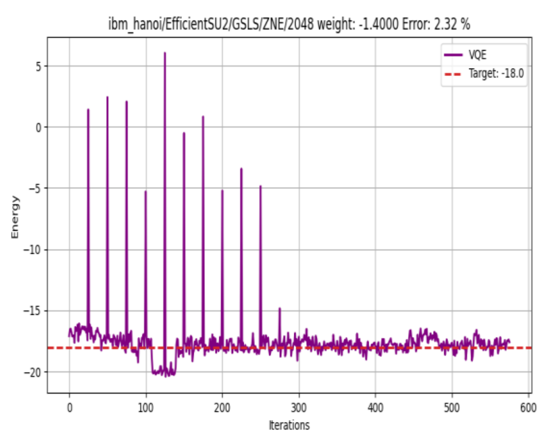

Figure 3 shows the shape of the curve using the efficient SU(2) Ansatz[2] made of two layers of single qubit SU(2) rotations and linear entanglements. It is coupled with the Gaussian-smoothed Line Search (GSLS)[3] optimizer based on Gaussian-smoothed samples of a sphere. The experiment ran for about 4h in the 27 qubit Hanoi processor.

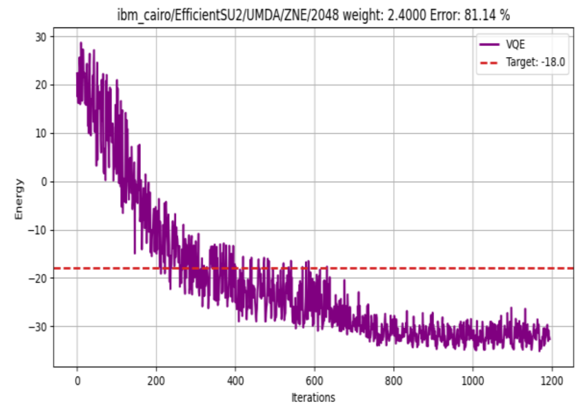

Figure 4 shows a bizarre shape with a very large number of cycles run in the 27 qubit Cairo processor using the Univariate Marginal Distribution Algorithm (UMDA), a stochastic search from the family of the evolutionary algorithms. [5]. This experiment took an excessive 5+h of run time.

After careful simulation our experiment settled with the following preparation:

1. **Ansatz of choice EfficientSU2:** It uses a well-known heuristic pattern to prepare trial wave functions common in classification tasks for machine learning (figure 5). For the sake of simplicity, and to avoid noise accrual, a single repetition of the gate sequence is used.
2. **Optimizer: NFT** - This is the Nakanishi-Fujii-Todo algorithm [6]. In our simulations, when combined with EfficientSU2, it produces the best curve of the lot: It descends quickly with a relatively low number of cycles (see figure 6).

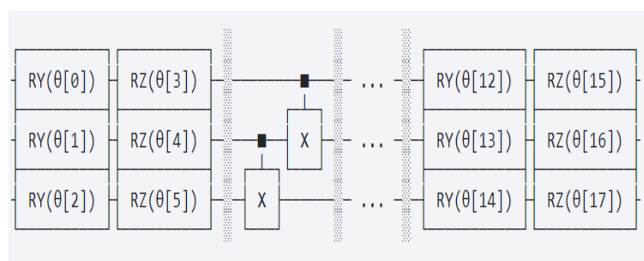

Figure 5 EfficientSU2 Ansatz for 3 qubits with linear entanglements. Note that the RY-RZ-CX pattern repeats an arbitrary number of times (implementation specific).

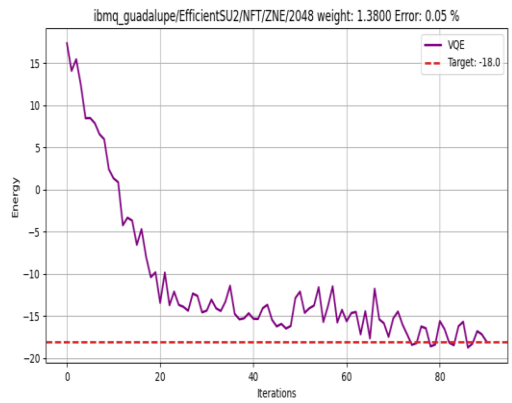

Figure 6 NFT optimizer combined with EfficientSU2 descends quickly to the ground state with a low number of cycles. This experiment was run on IBM's Guadalupe processor and produced a very high fidelity of 0.05 below 1%.

The Heisenberg spin ½ Hamiltonian is the final piece of the puzzle (see figure 8). It is made up of interaction of the Pauli Matrices X, Y, Z over neighbor qubits. It expands to 54 terms of observables each made of the product of a Uniform Interaction (initialized to the unit weight of the edge of the lattice) times a 16-tensor product of Identities (I) and Pauli Matrices mapped to the quantum processor qubit layout, producing a 2^{16} square density matrix. The Uniform interaction plays a critical role in error mitigation in our solution.

$$H = \sum_{<i,j>}^{N} X_i X_j + Y_i Y_j + Z_i Z_j \quad (1)$$

IV. In Noiseless Simulation, it runs like a Dream

Initial experiments in the noiseless simulator were very encouraging. We couldn't go wrong with such high fidelities, as they will counter even the noisiest of environments. Our assumption was wrong, the more we run on hardware, the more our joy turns into dejection.

V. On Hardware: Noise is King

We got lucky a few times on Hardware achieving high fidelity, however the more runs we tried, the affair started to look a lot like gambling at the casino: Get lucky a few times, but lose your shirt at the end. During the 4 months spent in this project, the noise accrued from two sources: entanglements (CX gates), and readouts (measurements). See figure 7.

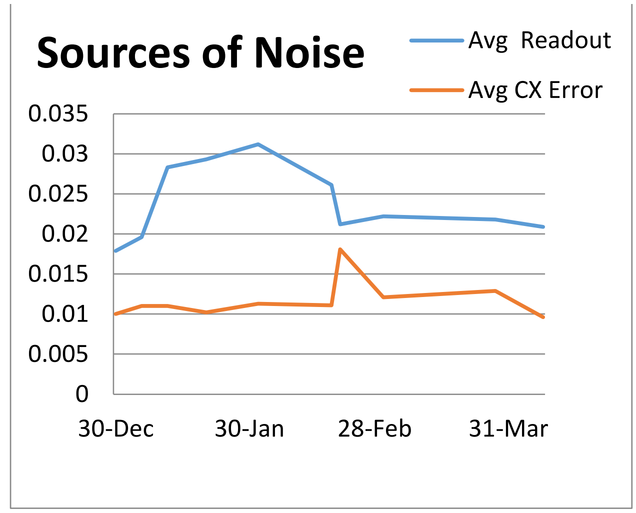

Figure 7 Averages for the two main sources of noise for the 4 month period of this experiment. Readout errors crept up around twice the level of CX gates.

Average noise levels appear low: between 1-3% in CX and readout. However, the amounts accrue with a large number of qubits (12 readouts in our case) and 11 CX gates for a 1 rep of the EfficientSU2 Ansatz. Wrangling the noise became the toughest part of the experiment.

VI. IQuantum Resilience: No Silver Bullet

Qiskit already features a sophisticated quantum resilience architecture [7] which was included in our initial design. However, it made little difference if any. Table 1 shows results for the three levels of resilience available: Twirled readout error extinction (T-Rex), Zero Noise Extrapolation (ZNE), and Probabilistic error cancellation (PEC).

Table 1: Quantum resilience failure rates for 114 experiments on stage 1.

Method	Failure	Success
T-Rex	20	2
ZNE	84	7
PEC	1	0

VII. Initial Experimental Results were very Noisy

We were disappointed when the final metrics were collected from hardware (see Table 2 Figure 8). Besides the high error level and disappointing fidelity, the noise tolerance was bad: From a total of 114 experiments, 104 failed. A great deal of effort was put on collecting this data, with an average wait time of 4 days in the execution queue and a staggering 170K jobs totaling 228h of quantum time.

Table 2 Metrics collected with Quantum resilience over a span of 4 months.

Total experiments	114
Failed	104
Avg # of VQE cycles/experiment	150
Total Jobs	17000
Avg execution time (h)	2
Total Quantum time (h)	228
Avg Queue wait time (days)	4

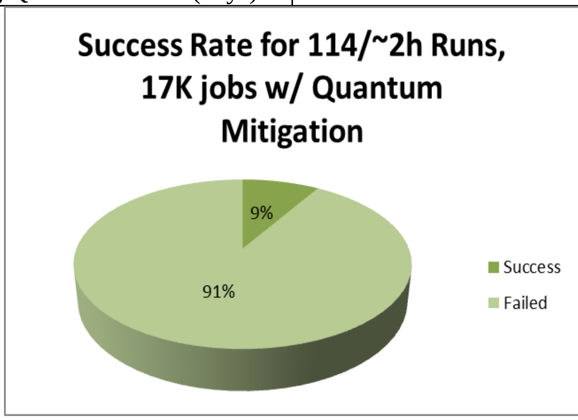

VIII. Error Mitigation by Classical Means?

We faced a difficult task: Find a way to reduce not only the noise of a single experiment, but build a solution that is noise tolerant over multiple runs. Error correction codes: such as Steane[8] or Shor[9] are out due to the high number of ancilla qubits required to stabilize a single physical qubit. Classical post processing techniques such as Fourier analysis were not feasible due to time constraints. The only choice available was to try some sort of algorithmic post processing: Instead of massaging the data, massage the logic that consumes it.

IX. Algorithmic Mitigation to the Rescue

Using tried and true Object Oriented Design Techniques, we tried a last ditch attempt to wrangle the noise. A fundamental Object Oriented design principle states that objects should be immutable (this means, they cannot be changed after allocation). This has the benefit of preventing mistakes in highly concurrent environments. Thus, in our solution, arguments sent to the VQE (Ansatz, Hamiltonian, and Optimizer) are immutable by default. However, for this situation, the Hamiltonian was turned into a mutable object within the VQE so it can be corrected dynamically by applying four simple algorithmic rules. To understand this, imagine a gambler at the casino. He can get:

1. Lucky: If a point falls below the desired fidelity or error threshold of 1%, abort the process and return the collected data.
2. Too lucky: If a point falls below the target ground state by some delta (the distance between the point in the curve and the classic ground state), dynamically decrease the uniform interaction (UI) of the Hamiltonian, then continue. This has the effect of driving the curve upwards to the ground state.
3. Unlucky: if the point falls above target by delta, do the inverse: Increase the UI & continue. This drives the curve downwards to the target ground state.
4. Recurse: If at the end of the optimization cycle, the point ends above the target, continue from the last point: The optimization resumes towards the ground state. As a failsafe to prevent infinite recursions, a max 5 recursive calls is allowed. If at the end, the ground state is not reached, the experiment fails.

These four rules work in tandem with the probability of rule one being triggered at any stage of the process. This seemed like an unorthodox solution; nevertheless it flipped the script 180 degrees! The results were pleasantly surprising (see Table 3 Figure 9).

Table 3 Final metrics collected with algorithmic mitigation.

Total experiments	76
Failed	1
Avg # of VQE cycles/experiment	100
Total Jobs	7600
Avg execution time (h)	1
Total Quantum time (h)	76
Avg Queue wait time (days)	6

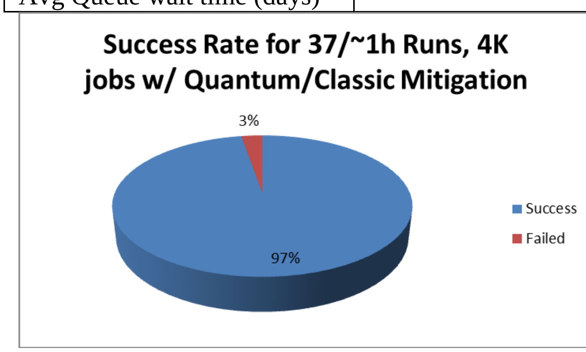

X. Conclusion

All in all, our solution for the Kagome lattice achieves high noise tolerance (it reaches a fidelity or relative error below 1% of the classical Eigen solver, 99% of the time). It does this using simple object oriented algorithmic post processing (as a matter of fact, the magic code that does the trick is less than 50 lines). When you run our experiment, it zigzags through the quantum noise with a high probability of success. Noise played a huge role in this project with the lion share of the time spent in error mitigation alone (see Table 4). It is critical to find a feasible solution to this problem. You don't want to spend 80% of your time dealing with errors.

Table 4 Time distribution (h) for this project over a period of 4 months.

Design	24 (6.7%)
Implementation	6 (1.7%)
Testing	12 (3,4%)
Documentation	4 (1.1%)
Noise Mitigation - Design	30 (8.4%)
Noise Mitigation - Implementation	280 (78.7%)

ACKNOWLEDGMENT

We would like to thank IBM for opening these incredible machines to the academic and scientific communities for experimentation; as well as the IBM Quantum Awards – Open Science Prize 2022 for which this work was made.

REFERENCES

[1] Balents, L. Spin liquids in frustrated magnets. Nature 464, 199–208 (2010). https://doi.org/10.1038/nature08917

[2] Qiskit Circuit Library available online at https://qiskit.org/documentation/stubs/qiskit.circuit.library.EfficientSU2.html.

[3] GSLS from Circuit Library available online at https://qiskit.org/documentation/stubs/qiskit.algorithms.optimizers.GSLS.html.

[4] McClean et al., Barren plateaus in quantum neural network training landscapes. arXiv:1803.11173.

[5] UMDA circuit library available online at https://qiskit.org/documentation/stubs/qiskit.algorithms.optimizers.UMDA.html.

[6] NFT optimizer available at https://arxiv.org/abs/1903.12166.

[7] Error suppression and error mitigation with Qiskit Runtime available at https://qiskit.org/documentation/partners/qiskit_ibm_runtime/tutorials/Error-Suppression-and-Error-Mitigation.html.

[8] Steane, Andrew (1996). "Multiple-Particle Interference and Quantum Error Correction". Proc. R. Soc. Lond. A. 452 (1954): 2551–2577. arXiv:quant-ph/9601029. Bibcode:1996RSPSA.452.2551S. doi:10.1098/rspa.1996.0136. S2CID 8246615.

[9] W.Shor, Peter (1995). "Scheme for reducing decoherence in quantum computer memory". Physical Review A. 52 (4): R2493–R2496. Bibcode:1995PhRvA..52.2493S. doi:10.1103/PhysRevA.52.R2493. PMID 9912632.

Network models on quantum photonic circuits

Author: Daniel Soh, Sandia National Laboratories, Livermore, California (dbsoh@sandia.gov)
Topic: Computing and programming models – quantum network models and architecture

Challenge

Today's physics-level models for describing coherently connected quantum systems face significant challenges as quantum network technology advances rapidly towards modular, integrated circuits. The quantum network community has adopted the Scattering, Lindbladian, and Hamiltonian (SLH) model, originating from Gardiner and Collet's open quantum system model [1] and developed by James and Gugh [2], to simplify the complex quantum networks comprising multiple interconnected quantum systems. While the SLH formalism has successfully described and predicted the performance of various quantum network systems (primarily free-space connected quantum systems), it lacks several essential features that are particularly relevant to integrated photonic quantum circuits. These shortcomings make it difficult to apply the existing SLH formalism to state-of-the-art connected quantum systems on photonic chips. The missing features include appropriate handling of time delays and dispersion in connection lines. However, the most critical aspect that the SLH formalism currently lacks is the treatment of non-Markovian baths, where the environment's timescale is not much faster than the primary quantum system's, rendering the memory-less bath assumption inapplicable [3]. As we increasingly move toward integrated photonic and phononic circuits [4], addressing the small bath and its associated memory effects becomes increasingly crucial. At present, our understanding of an SLH-like formalism for non-Markovian baths remains limited.

Opportunity

To address the diverse needs of quantum network modeling, it is crucial to incorporate features such as time delays, dispersion, and non-Markovian baths while deriving network reduction rules like concatenation, series, and feedback. The integration of time delays and non-Markovianity can make the feedback reduction rule especially challenging due to the complex dynamics involving memory effects. Nevertheless, several classical control theory approaches show promise in tackling these issues.

One such approach is the receding-horizon framework, which gradually truncates the memory effect over time [5]. Additionally, quantum model reduction techniques can be employed, where mean-field approximations replace average dynamic behaviors while preserving the most prominent quantum features through perturbative quantum noise [6]. Combining these techniques with the fundamental treatment of time delays using cascaded cavities [7] and the direct derivation of non-Markovian connected quantum systems [8] may lead to a comprehensive approach for practical connected quantum systems. As a result, the reduced quantum system could effectively predict the dynamics and performance of quantum networks.

Assessment

The new SLH-like framework, which now integrates time delay, dispersion, and non-Markovian baths, can be experimentally tested. For instance, Fig. 1 presents a schematic of a feedback network consisting of an optical parametric amplifier light squeezer and a dispersive time-delay element. By measuring the squeezing spectrum of the output field, one can compare the experimental results with predictions from the enhanced model based on this novel SLH-like framework.

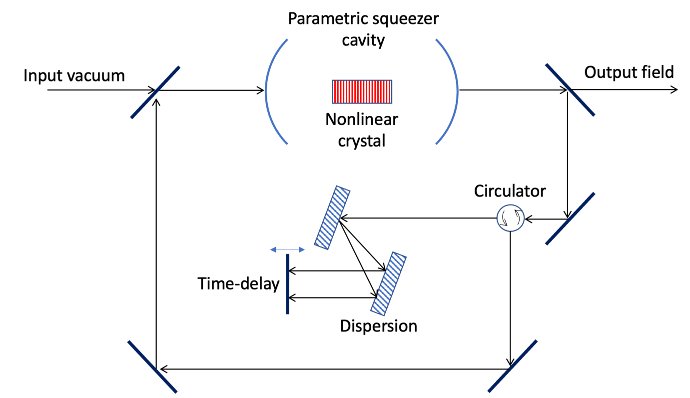

Fig. 1. Schematic of a test experimental setup to examine the effect of a dispersive connecting bath with a time delay on the output squeezing spectrum.

To test the non-Markovian bath, an integrated photonic circuit can be utilized. In this circuit, the intermediate waveguide (bath) connects two primary quantum photonic systems (e.g., cavities) and is generally short, although its length significantly exceeds the light's wavelength. Consequently, it is anticipated that the bath in typical photonic circuits may display non-Markovian features, such as the memory effect. One can fabricate a photonic circuit resembling the schematic depicted in Fig. 1. Comparing the experimental results with predictions from the enhanced model will help evaluate the efficacy of the new SLH-like quantum network framework.

Timeliness or maturity

The significant progress in integrated quantum photonic circuits has enabled coherent connections between quantum systems, allowing them to achieve complex quantum dynamics and functions that were previously unattainable in single quantum systems. In particular, we are now witnessing large-scale quantum networks with multiple connected qubit memories and gates. This recent advancement highlights the need to develop new network modeling capabilities that incorporate previously neglected features in existing network theories.

Moreover, the rapid emergence of new insights into non-Markovian systems and advancements in model reduction techniques make it a timely endeavor to address the challenging problem of non-Markovian quantum networks with dispersion and time delays.

References

[1] Gardiner, Crispin W., and Matthew J. Collett. "Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation." *Physical Review A* 31.6 (1985): 3761.
[2] Gough, J. J. M. R., and Matthew R. James. "Quantum feedback networks: Hamiltonian formulation." *Communications in Mathematical Physics* 287 (2009): 1109-1132.
[3] Combes, Joshua, Joseph Kerckhoff, and Mohan Sarovar. "The SLH framework for modeling quantum input-output networks." *Advances in Physics: X* 2.3 (2017): 784-888.
[4] Elshaari, Ali W., et al. "Hybrid integrated quantum photonic circuits." *Nature Photonics* 14.5 (2020): 285-298.
[5] Michalska, Hannah, and David Q. Mayne. "Robust receding horizon control of constrained nonlinear systems." *IEEE transactions on automatic control* 38.11 (1993): 1623-1633.
[6] Tezak, Nikolas, Nina H. Amini, and Hideo Mabuchi. "Low-dimensional manifolds for exact representation of open quantum systems." *Physical Review A* 96.6 (2017): 062113.
[7] Tabak, Gil, and Hideo Mabuchi. "Trapped modes in linear quantum stochastic networks with delays." *EPJ Quantum Technology* 3.1 (2016): 3.
[8] Zhang, Wei-Min, et al. "General non-Markovian dynamics of open quantum systems." *Physical review letters* 109.17 (2012): 170402.

Quantum Algorithms for Characterizing Experimental Quantum Systems

Authors: Andrew Sornborger (LANL, sornborg@lanl.gov, presenter), Alberto Marino (ORNL), Claire Marvinney (ORNL), Tyler Volkoff (LANL), Andrea Delgado (ORNL), Zoë Holmes (EPFL), Raymond Newell (LANL), Mike Martin (LANL), Malcolm Boshier (LANL), Michael McKerns (LANL)

Topic: Applications, Algorithms, Quantum Compilation, Co-Design

Challenge: Standard (classical) methodology for characterizing quantum systems can be cumbersome and inefficient. With classical methodology, characterizing an experiment requires construction and testing of a system model that often takes years of effort. Methods such as spectral estimation, process tomography, and scattering (among many others) are all used, potentially in combination with a feedback loop between experiment and theory, to determine the precise properties of the system. The exponential scaling of these techniques with the size of the quantum system makes it challenging to implement them for realistic systems of interest for quantum computing and quantum networking. It is therefore incumbent on us to consider, design and implement new quantum methods with better scaling for full characterization of complex laboratory-scale quantum systems if we wish to understand the capabilities and limitations of current quantum technology and prepare for a future of fully connected quantum networks of quantum computers.

Opportunity: With the advent of new concepts from the field of quantum information science, investigators have begun to study whether quantum algorithms are useful to accelerate our capabilities for system characterization. Crucially, in contrast to the *classical* descriptions provided by conventional tomographic methods, quantum algorithms can be used to learn *quantum* (i.e. circuit based) descriptions of experimental processes (see Fig. 1). This shift in approach provides a means of sidestepping the exponential complexity of standard process tomography. Key to this approach is the observation that quantum computers, as inherently quantum systems, are much more naturally suited to modelling other quantum systems. New advances in quantum machine learning (QML) have further shown that entanglement, a natural resource in many photonic experiments, can be used to reduce the amount of training data needed to learn the full dynamics of a quantum system. In particular, quantum entanglement can be used not only to model the target quantum process but also as a resource to enhance the learning efficiency. These approaches have not yet reached the implementation stage and are only being currently considered for simple (few parameter) quantum systems; however, they could be very valuable in order to push quantum technologies such as quantum algorithms for sensing and computation forward in laboratory-scale and even larger, networked, experiments.

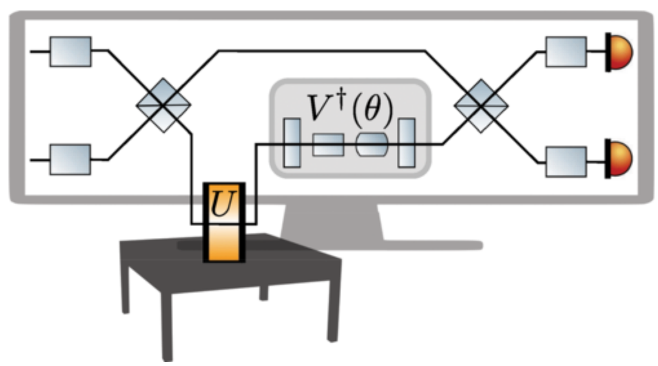

Fig. 1: A QML algorithm for characterizing a quantum system. QML is used to learn a parametrized quantum model, $V(\theta)$, of an experimental photonic system, U.

These novel quantum system characterization approaches are compatible with different quantum systems. Among them, continuous variable (CV) quantum optical systems, a primary candidate for quantum computation, provide a path forward as they provide deterministic operation, scalability, and quantum error mitigation schemes [1, 2]. The novel QML methods considered here have been shown theoretically to offer an advantage for quantum system characterization of complex

CV quantum systems [3]. Such methods open up new approaches to characterizing the properties (e.g. the processes implemented) of novel optical materials. More generally, these algorithms could be used as a subroutine in more complex algorithms intended to learn effective Hamiltonians to understand the emergent properties of quantum materials. Furthermore, the resources needed to scale up quantum system characterization to more complex systems, such as CV quantum optical systems, are also of use for quantum networks and quantum computing, thus offering opportunities for co-design toward larger, fully connected, quantum networks of sensors and computers.

A second platform where QML-based characterization algorithms could be of use is neutral atom-based quantum simulators, of interest for their long coherence times [4]. Here, smaller systems of entangled probe atoms could interact with a large, quantum many-body system of interest. Extracting useful information from this interaction requires extending the technique to characterizing a quantum channel [5], an important task for any generalized quantum network with loss or decoherence. The unifying feature is the connection of many-body dynamics and emergent phenomena (an analog quantum computing task), with a lossy quantum channel and techniques for its characterization.

Assessment: The goal of QML-based characterization is to learn a quantum model formulated in a quantum circuit that could be used to fully reproduce the behavior of an experimental quantum system. Performance may be evaluated by direct comparison of the learned quantum model with the experimental system via a cost encoded in the learning algorithm. Metrics of success would be 1) the ability to efficiently minimize an experimentally realized cost function to evaluate the fidelity (or other appropriate measure) between the experiment and model and 2) the demonstration of the efficiency (i.e. polynomial scaling with the size of the quantum system) of quantum-enhanced characterization methods (in contrast to exponential scaling of conventional tomography). How well entanglement can be realized must be assessed for the quantum system used to implement the algorithm if entanglement is used to compute the cost. The successful extension to configurations with more entangled modes/systems for characterization of a complex quantum system will also serve as a metric of success, as well as the scaling with system complexity of the required quantum resources.

Timeliness or maturity: The congruence of recent progress in quantum machine learning with advances in entanglement generation in experimental quantum hardware gives us a realistic and timely opportunity to realize an exponential improvement in the efficiency of characterization of complex quantum systems. The success of this effort will accelerate progress on quantum computing platforms and the deployment of quantum networks and quantum computing hardware.

References:

[1] S. L. Braunstein and P. van Loock, “Quantum Information with Continuous Variables,” *Review of Modern Physics* **77**, 513 (2005).

[2] D. Su, R. Israel, K. Sharma, H. Qi, I. Dhand, and K. Brádler, “Error Mitigation on a Near-Term Quantum Photonic Device,” *Quantum* **5**, 452 (2021).

[3] T. Volkoff, Z. Holmes, and A. Sornborger, “Universal Compiling and (No-)Free-Lunch Theorems for Continuous Variable Quantum Learning,” *PRX Quantum* **2**, 040327 (2021).

[4] G. Giudici, M. D. Lukin, and H. Pichler. "Dynamical Preparation of Quantum Spin Liquids in Rydberg Atom Arrays," *Physical Review Letters* **129**, 090401 (2022).

[5] N. Ezzell, et al. "Quantum Mixed State Compiling," *Quantum Science and Technology* **8**, 035001 (2023).

Buffered Entangled Pairs for Asynchronous Quantum Networking

Samuel A. Stein[1], Michael Austin DeMarco[2,3], Chenxu Liu[1], James Ang[1], and Ang Li[1]

[1]Pacific Northwest National Laboratory
[2]Massachusetts Institute of Technology
[3]Brookhaven National Laboratory

Topic

Error mitigation in Networking Applications

Challenge

Quantum networking hinges on the capacity to enable distributed quantum communication between various systems, spanning from quantum processors within a single dilution refrigerator to those separated by hundreds of miles. The precise method for achieving efficient and concise communication between processors remains an unresolved issue, although distributing maximally entangled states between nodes appears to be a promising strategy for facilitating communication. One method of communication is distributed entanglement for gate teleportation, a practical approach for realizing distributed communication between processors. In particular, gate-communication protocols that utilize Bell state consumption in conjunction with classical communication enable the execution of controlled operations between two physically distant systems.

The efficacy of these communications largely depends on the availability of a high-purity distributed Bell state. These states are generated usually via a probabilistic entanglement generation process. Executing remote quantum gate operations over vast distances, with multiple possible connection paths, on a microsecond timescale, presents a considerable challenge. Establishing a quantum network infrastructure necessitates the creation of a time-sensitive resource that ensures consistent high fidelity, adapts to varying demand and production, and navigates a complex web of connections. Achieving this while maintaining practical hardware objectives poses a formidable task.

Opportunity

Entanglement Distillation (1–4) facilitates the purification of low-fidelity entangled pairs into a reduced set of higher-fidelity pairs. In theory, perfect gates and long-lived quantum states would enable purification to reach unit fidelity. However, the practical implementation of these protocols encounters various obstacles, such as the generation rates of corresponding Bell pairs, initial Bell pair fidelity, and idle errors during the wait for a communication request. In a recent paper (5), photon heralded entanglement generation schemes are considered. With the modeling of microwave-to-optical transduction process, the Bell states fidelities and generation rates are calculated. As illustrated in Figure 1, the raw Bell pairs' inadequate fidelities for distributed communications are brought to light. Moreover, the probabilistic nature of the entanglement generation scheme underscores a primary challenge in the physical implementation of entanglement distillation protocols: the task shifts from purifying a large set of low-purity entangled pairs into a smaller, high-purity set, to continuously distilling a stream of low-purity entangled pairs while managing their demand and striking a balance between load and performance.

Regarding distillation performance, as depicted in Figure 2, the ideal distillation process necessitates two pairs of the same fidelity, resulting in an exponential distillation cost since each purification round consumes two pairs to produce one (1). Nonetheless, when noise is present, pairs that decohere to lower fidelities can be combined with earlier purification rounds to maintain a consistent fidelity level. This can be conceptualized as distilling a Round 3 pair (Fidelity $F = 0.98$) with a Round 6 pair ($F = 0.98$) to generate a Round 4 pair ($F = 0.99$).

In the past, superconducting systems featuring transmons excelled in high gate fidelities and fast gate execution. However, transmons' lifetimes are relatively short, ranging from hundreds of microseconds to a potential future span of a few milliseconds (6, 7). This limitation poses a significant challenge for networking, as it prevents Bell pairs from being buffered during communication between nodes and hinders the creation of a stable connected quantum network.

Recently, multimode cavities paired with SNAILs (8) and transmons have demonstrated not only exceptionally long lifespans on the order of tens of milliseconds but also the capability to store multiple qubits within a single cavity (9). Furthermore, using quantum error correction codes to encode the microwave cavity field, the lifetime can be further extended (10). This development presents a valuable opportunity for communication buffering.

The recent proposal of HetArch (11) introduces a toolbox for designing Quantum Heterogeneous Microarchitectures. It suggests the use of chiplets tailored for specific tasks, including a sample entanglement distillation module, which relies heavily on these multimode cavities. Adopting a similar approach for quantum networks could yield significant benefits. In terms of buffering Bell pairs, a microarchitecture akin to HetArch's Entanglement Distillation Controller could be employed, incorporating alternative micro-controller logic.

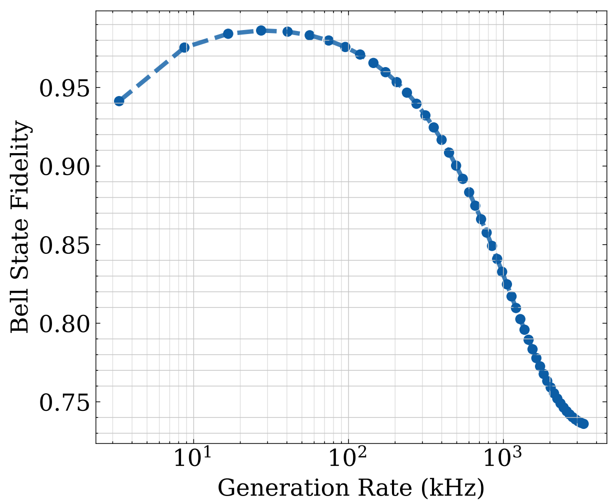

Fig. 1. Raw Entangled Pair Fidelity as a function of Generation Rate over Spontaneous Parametric Down Conversion, a quantum optics entangling process

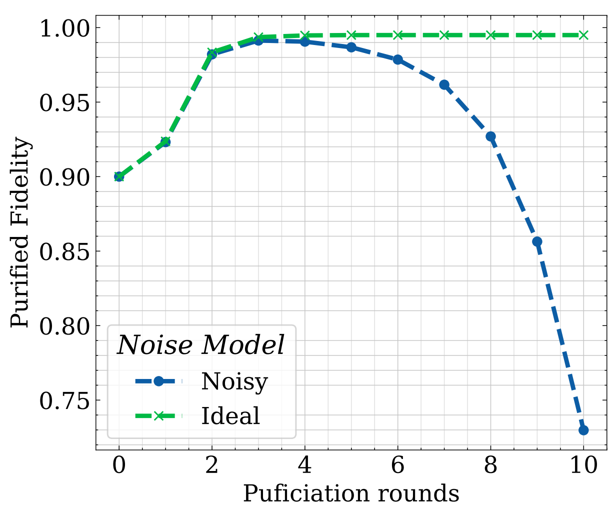

Fig. 2. Distillation Performance of the Deutsch-Josza distillation protocol (2) based on an initial Werner State fidelity of 0.90 and generation time of 1 µs.

While HetArch's controller distills up to a target point and outputs the pair, an opportunity arises to reconfigure the controller logic to accommodate the variable demands of quantum networks.

To ensure high-fidelity connections between each node in quantum networks, nodes should be capable of requesting Bell pairs on demand. Prolonged wait times for communication can lead to decoherence and idling errors, burdening local compute nodes along an avoidable axis. Such issues can be mitigated through more effective network engineering. The challenge lies in harmonizing network demand with distilled Bell pair generation, understanding the role of distillation in network scheduling, and addressing the myriad design considerations when constructing large-scale networks.

Assessment

To assess the performance of how well we can buffer Bell pairs and perform remote communication relies on understanding the limitations of buffering and distillation over a network. Increased distillation reduces the potential buffered Bell states availability at a given moment, however provides increased communication fidelity. Assessing a quantum network will require abstractions of system layers. Performance analysis of the local buffering controller, local network topology, and global network topology requires investigation. Latency and fidelity characterize each layers performance. Increasing fidelity comes at the cost of increasing latency, as lower generation rates and increased number of distillation rounds both can increase fidelity, however at the cost of time.

Timeliness or maturity

Recent key developments in superconducting quantum computing, specifically multi-modal cavities, have allowed for relatively long lifetimes, a key requirement for buffering. Although the long lifetime is insufficient to guarantee long-lasting distributed quantum states across multiple quantum nodes, it enables entanglement purification and other error correction and mitigation techniques. Furthermore, Heterogeneous Quantum Computing has recently sparked much interest due to technology-specific advantages, with HetArch proposing an Entanglement Distillation Module comprising multi-mode resonator memories coupled to transmons' to generate high fidelity Bell states.

References

1. Charles H. Bennett, Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A. Smolin, and William K. Wootters. Purification of noisy entanglement and faithful teleportation via noisy channels. *Phys. Rev. Lett.*, 76:722–725, Jan 1996. doi: 10.1103/PhysRevLett.76.722.
2. David Deutsch, Artur Ekert, Richard Jozsa, Chiara Macchiavello, Sandu Popescu, and Anna Sanpera. Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Physical review letters*, 77(13):2818, 1996.
3. W Dür and H J Briegel. Entanglement purification and quantum error correction. *Reports on Progress in Physics*, 70(8):1381, jul 2007. doi: 10.1088/0034-4885/70/8/R03.
4. Stefan Krastanov, Victor V Albert, and Liang Jiang. Optimized entanglement purification. *Quantum*, 3:123, 2019.
5. James Ang, Gabriella Carini, Yanzhu Chen, Isaac Chuang, Michael Austin DeMarco, Sophia E Economou, Alec Eickbusch, Andrei Faraon, Kai-Mei Fu, Steven M Girvin, et al. Architectures for multinode superconducting quantum computers. *arXiv preprint arXiv:2212.06167*, 2022.
6. Alexander P. M. Place, Lila V. H. Rodgers, Pranav Mundada, Basil M. Smitham, Mattias Fitzpatrick, Zhaoqi Leng, Anjali Premkumar, Jacob Bryon, Andrei Vrajitoarea, Sara Sussman, Guangming Cheng, Trisha Madhavan, Harshvardhan K. Babla, Xuan Hoang Le, Youqi Gang, Berthold Jäck, András Gyenis, Nan Yao, Robert J. Cava, Nathalie P. de Leon, and Andrew A. Houck. New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *Nature Communications*, 12(1):1779, Mar 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-22030-5.
7. Chenlu Wang, Xuegang Li, Huikai Xu, Zhiyuan Li, Junhua Wang, Zhen Yang, Zhenyu Mi, Xuehui Liang, Tang Su, Chuhong Yang, Guangyue Wang, Wenyan Wang, Yongchao Li, Mo Chen, Chengyao Li, Kehuan Linghu, Jiaxiu Han, Yingshan Zhang, Yulong Feng, Yu Song, Teng Ma, Jingning Zhang, Ruixia Wang, Peng Zhao, Weiyang Liu, Guangming Xue, Yirong Jin, and Haifeng Yu. Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds. *npj Quantum Information*, 8(1):3, Jan 2022. ISSN 2056-6387. doi: 10.1038/s41534-021-00510-2.
8. N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret. 3-wave mixing Josephson dipole element. *Applied Physics Letters*, 110(22), 05 2017. ISSN 0003-6951. doi: 10.1063/1.4984142. 222603.
9. Srivatsan Chakram, Andrew E Oriani, Ravi K Naik, Akash V Dixit, Kevin He, Ankur Agrawal, Hyeokshin Kwon, and David I Schuster. Seamless high-q microwave cavities for multimode circuit quantum electrodynamics. *Physical review letters*, 127(10):107701, 2021.
10. V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret. Real-time quantum error correction beyond break-even. *Nature*, 616(7955):50–55, Apr 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-05782-6.
11. Samuel Stein, Sara Sussman, Teague Tomesh, Charles Guinn, Esin Tureci, Sophia F Lin, Wei Tang, James Ang, Srivatsan Chakram, Ang Li, Margaret Martonosi, Fred Chong, Andrew A. Houck, Isaac L. Chuang, and Michael DeMarco. Hetarch: Heterogeneous microarchitectures forsuperconducting quantum systems. *arXiv preprint To Be Released*, 2023.

An Assessment of Quantum Machine Learning Challenges and Opportunities

Omer Subasi and Joseph Manzano
omer.subasi@pnnl.gov - joseph.manzano@pnnl.gov
Pacific Northwest National Laboratory (PNNL), Richland, Washington, USA

In this position paper, we discuss the challenges and opportunities for **Quantum Machine Learning (QML)**, which is categorized in **Quantum Algorithms**. We populate six main challenges and organize our work accordingly (Figure 1).

A. Challenges

1- Quantum advantages: There are significant doubts regarding the reported quantum advantages in the QML literature [1]. The doubts concentrate around if (hybrid) QML can be as successful as the modern industrial-scale classical ML, such as deep learning (DL) [2]. Such claims are typically achieved under non-practical settings or assumptions. Moreover, the success stories with small-scale datasets and simulations are not sufficient, although, at the moment, it is not possible to be test or simulate large-scale QML.

2 - The barren plateaus: As the number of qubits and the size of quantum parameterized circuits increase, novel problems occur in the loss function landscapes. Barren plateaus [3] are exponentially flat landscapes in the number of qubits or circuit size/depth. That is, the gradients vanish exponentially as the number of qubits and circuit sizes increase. This makes finding a global minimum exponentially expensive in terms of the number of shots (or the resolution) required to make progress in the optimization. As a result, barren plateaus seem to nullify potential quantum advantages.

3 - The impact of quantum noise: Currently available and near-term quantum devices are and will be noisy [4]. Existing QML studies mostly ignore the impact of noise. However, it is reported that noise can induce the barren plateau problem [5], and corrupt the loss landscape by exponentially suppressing or erasing the relevant features [6].

4 - The lack of large-scale fully quantum datasets: As the latest research seems to suggest that achieving a quantum advantage in QML with classical data is very hard. Alternatively, QML with fully quantum data might enable an advantage relatively easier. However, there exists only a handful of small-sized fully quantum datasets [7], [8]. Large-scale fully quantum data is missing.

5 - Data encoding and feature maps: Efficient data encoding methods, i.e., feature maps, are crucial for processing classical data on quantum computers as well as the success of variational quantum learning [9]. Developing high-performing encoding methods remains an open problem.

6 - Theoretical understanding of QML: Gaining a deeper theoretical understanding of the QML fundamental principles and limitations is vital for guiding the development of new algorithms. Theoretical studies typically do not get the appropriate investment due to the prioritization of deliverable products and short-term goals. Moreover, the difficulty and uncertainty about the progress/gain/return with the investments seem to discourage theory-oriented studies.

B. Opportunities

1 - QML foundations: As it is not possible to obtain and test a quantum advantage at scale right now, some research focus can instead be developing the building blocks for QML similar to the perceptron of classical neural networks. It may also be beneficial to investigate the other aspects in QML, such as different cost functions, architectures, and types of feature reduction methods.

2 - New algorithms and efficient co-design: Theoretical studies on barren plateaus, their causes and mitigation are urgently needed. In addition, understanding the computational complexities and power of hybrid quantum and classical QML can shed light on the performance scaling and its feasibility. Moreover, efficient integration between hybrid quantum and classical hardware can also targeted.

3 - Novel noise-aware technologies and algorithms: Noise-resistant QML algorithms are needed to be developed. At the very least, studies need to evaluate and report the impact of quantum noise on the performance of their solutions. Hardware providers can enhance their technologies to limit the effects of quantum noise.

4 - Fully Quantum Datasets: Large-scale fully quantum datasets and tools need to be developed.

5 - New data encoding methods and technologies: New encoding methods for classical data may unlock quantum advantages if the methods are linear or sub-linear. To implement such methods, new hardware and software technologies will have to follow.

6 - Advanced theories and collaborations in research: New theories explaining the expressive power of QML methods as well as their computational complexities will help drive new algorithmic solutions to the existing problems such as performance scaling, the need for efficient algorithm designs, and the seemingly road-blockers such as barren plateaus. At the most general settings, a unified theory of QML will help us understand the fundamental and common properties and principles across different QML algorithms and approaches. It will also help us to incorporate our prior knowledge into

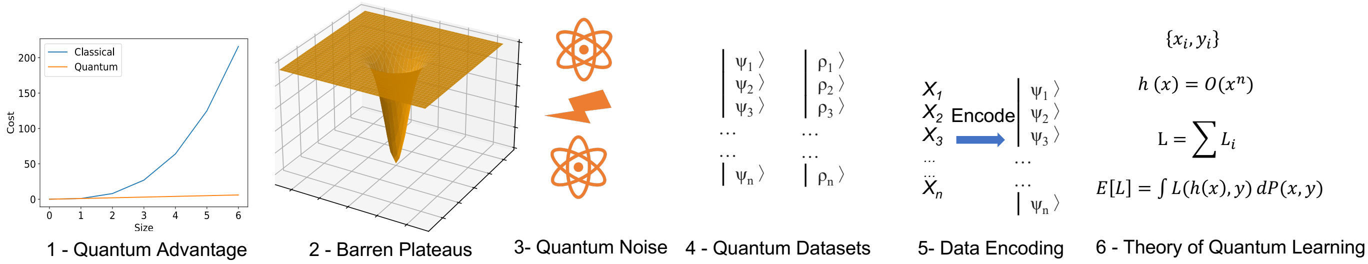

Fig. 1. QML Challenges

new solutions. Classical DL has already got powerful unifying theories such as [10].

C. Assessment

1- The success for this challenge would result in QML architectures (circuits) that are robustly trainable and scalable. The performance of these architectures would show verifiable quantum advantages in terms of metrics such as asymptotic run time. Moreover, a (preferably simple) building block, similar to the classical perceptron in neural networks, or self-attention in transformers, would be attained and employed across various architectures and/or learning tasks.

2- The barren plateau problem would be mitigated. This would mean even at large scales, QML architectures would be trainable. Hybrid QML would scale well and be efficient.

3- The long-term success for hardware technologies in this challenge would be quantum computers that are fault-tolerant just as the commercially available personal computers, those in data centers or supercomputers in national laboratories. The midterm success would be quantum devices that have high-fidelity circuits, and experience uncorrected noise / errors less frequently than today's quantum devices do. In terms of software and algorithms, uncorrected errors would not always lead to corrupt data or program crash since there would be software-based mitigation to counter the errors.

4- Reproducible large-scale training and testing based on fully quantum data would be available. We would be able to test and validate quantum advantages with real-world quantum data. The QML progress in terms of new applications would be faster and scalable.

5- Provided that efficient data encoding techniques are developed, hybrid quantum and classical QML would exhibit quantum advantages over classical ML.

6- A unified mathematical framework would be available to explain why QML methods and algorithms work the way they do or the way they perform. Such a framework would guide new designs and solutions that achieve a quantum advantage.

D. Timeliness

Quantum computing hardware and technologies are constantly improving. As a result, QML has just recently become testable and simulatable at small-scales. Similarly, quantum noise in these technologies has started to get much needed attention. With our first exposure to hybrid quantum and classical QML, researchers and scientist are converging to a consensus that a quantum advantage in hybrid QML seems to be very hard, if not infeasible. Consequently, the need to build fully quantum datasets has become urgent. Moreover, barren plateaus have been discovered in 2018 - relatively recently. We are at the conjunction that the existing quantum computing hardware enables meaningful research and deep learning, with its immense success, has yielded critical algorithmic and theoretical advances, and unforeseen approaches and applications. Now is the time to leverage relevant applicable ideas for QML.

If the success of deep learning can be repeated in QML, then it is very difficult to even imagine or predict what we could do. Likely, we do not know about some capabilities that we would obtain with large-scale trainable QML.

E. Concluding Remarks

We presented our assessment of the challenges and opportunities for QML. We posit that this is the right time to invest in QML to explore its potential and to see if it can achieve the immense success of modern classical ML and DL.

References

[1] Maria Schuld and Nathan Killoran. Is quantum advantage the right goal for quantum machine learning? *Prx Quantum*, 3(3):030101, 2022.

[2] Shi Dong, Ping Wang, and Khushnood Abbas. A survey on deep learning and its applications. *Computer Science Review*, 40:100379, 2021.

[3] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):4812, 2018.

[4] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermanni Heimonen, Jakob S Kottmann, Tim Menke, et al. Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1):015004, 2022.

[5] Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nature communications*, 12(1):6961, 2021.

[6] M Cerezo, Guillaume Verdon, Hsin-Yuan Huang, Lukasz Cincio, and Patrick J Coles. Challenges and opportunities in quantum machine learning. *Nature Computational Science*, 2(9):567–576, 2022.

[7] Louis Schatzki, Andrew Arrasmith, Patrick J Coles, and Marco Cerezo. Entangled datasets for quantum machine learning. *arXiv preprint arXiv:2109.03400*, 2021.

[8] Elija Perrier, Akram Youssry, and Chris Ferrie. Qdataset, quantum datasets for machine learning. *Scientific Data*, 9(1):582, 2022.

[9] Manuela Weigold, Johanna Barzen, Frank Leymann, and Marie Salm. Data encoding patterns for quantum computing. In *Proceedings of the 27th Conference on Pattern Languages of Programs*, 2022.

[10] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Fitting big data onto early quantum computers **Topic: algorithms**

Christoph Sünderhauf, Riverlane ⟨christoph.sunderhauf@riverlane.com⟩ April 2023

Challenge Quantum computing presents the alluring prospect of an exponential reduction in resources in application areas as diverse as computational chemistry or linear algebra. Crucially, this is enabled by the exponential state space of many-body quantum systems: Merely n qubits can explore and encode an exponential number $N = 2^n$ of data items. Yet, any exponential advantage may shrink away in the face of data loading. Loading N data items (eg. for preparing a general n-qubit state with N amplitudes) generally requires at least $O(2^n)$ gates and/or ancilla qubits, annihilating any exponential advantage.

To achieve favorable run-times with an exponential advantage, quantum algorithm proposals sometimes invoke QRAM [1], which has a small circuit depth, i.e. run-time $O(n)$. However, QRAM requires an exponential number of gates to be executed simultaneously on exponentially many ancillary qubits; the exponential resource requirement is in space rather than time. Alternatively, QROM [2] does not have an exponential ancillary qubit requirement, but an exponential circuit depth. The select-swap network [3] allows an interpolation between these regimes.

State preparation can be based on the data loading oracles mentioned above, or use other circuits implicitly loading the data; those methods also incur exponential space-time complexity [4]. Trotterization can be viewed as another implicit form of data loading; the cost increases with the number of data items (coefficients in the Hamiltonian).

Strategies to reduce data loading are *compressing* and *approximating* the data:

Truncating a basis expansion Truncating small amplitudes in the computational basis results in small data items close to 0 being ignored [5]. The Walsh basis is also cheap to prepare on a quantum computer and can be used [6].

Factorizing and truncating In quantum chemistry, a popular route to compress the data has been factorizing the Hamiltonian in different tensor factorizations, and then truncating an intermediate rank (see single factorization [5], double factorization [7], tensor hypercontraction factorization [8]).

Polynomial approximation If the data follows an analytic function, it can be compressed by approximating with a low-degree polynomial. Instead of loading the data, the polynomial can be implemented with QSVT [9].

Compressing repeated values If the data has a symmetry or otherwise contains values that are repeated or related to each other in a structured fashion, they only need to be loaded once [10], [11].

Neglecting small gates In a quantum circuit encoding the full data, rotation gates by small angles that are close to the identity are removed [12].

Statistical loading Only part of the data is loaded in each shot of the circuit [13], [14].

Opportunity New research directions to reduce the overall cost of data loading include:

Other truncation bases Other bases like the Fourier basis are efficient to prepare on QPUs and could result in stronger truncation.

Matrix-product states Tensor networks like matrix product states provide an efficient description for entangled many-body quantum states. Efficient data loading circuits based on these techniques could be developed.

Classical shadows Data readout has been accelerated by the powerful technique of classical shadows [15]. Related ideas for the inverse problem of data loading could be explored.

Classical-quantum changeover In hybrid quantum-classical algorithms, data readout and data loading occurs at every changeover between classical and quantum information processing. The overall data loading could be reduced by moving more computation from CPU to QPU.

For all directions, collaborations with experts from the subject areas could be fruitful. New software tools that automatically transform input data to an efficient quantum circuit selecting various of the above strategies could also be useful for the quantum computing community.

Assessment Reduced data loading should be evaluated from several perspectives:

Classical compute Preprocessing (like factorizations) and circuit synthesis can add significant classical cost.

Data loading oracle circuit length The QPU run-time for a single data loading oracle.

Repetitions of data loading oracle Due to the no-cloning theorem, if access to the data is required multiple times in a single circuit, the data loading oracle must be repeated. For example, in quantum chemistry the Hamiltonian is required in every iteration of quantum phase estimation, or in solving linear systems, the matrix to be inverted is required in multiple steps. Then, data loading is not merely an additive contribution to the algorithm's complexity, but multiplicative, directly affecting the asymptotic cost.

Number of circuit runs Hybrid quantum-classical algorithms require multiple circuits, increasing the overall data loading throughout execution of the algorithm.

Reasonable qubit requirement Ideally, circuits should be able to run on early fault-tolerant quantum computers with a limited number of logical qubits.

Timeliness or maturity In the quest for exponential speedup in applications based on linear algebra tasks, QRAM is often invoked as a necessary means to avoid the exponential run-time associated with data loading. However, its exponential ancillary qubit requirement makes it unfeasible for early fault-tolerant quantum computing devices with limited qubit count. As such, it is imperative to further study data loading if an exponential speedup is hoped to persist in practice.

Dequantizations results [16] promise a classical exponential speed-up and rely on a non-standard classical data access model, arguing that quantum advantage is related to quantum access to the data. Better understanding and improvements in quantum data loading is paramount to understand any room for quantum advantage [17].

Data loading's sister problem is called data readout. If all amplitudes of an output state are required, exponentially expensive full-state tomography must be carried out. Recent breakthroughs on classical shadows [15] indicate that for many observables, the cost can be reduced tremendously.

Data loading is a bottleneck even in quantum chemistry. For n orbitals, quantum computing allows to perform the ground state energy calculation in the full 2^n dimensional Hilbert space. While only a polynomial number of data items $O(n^4)$ is required to specify the Hamiltonian, the data loading is still a dominant contribution to the overall cost of the algorithm. Recent breakthroughs in quantum chemistry are due to reduction in data loading, both in oracle circuit length and repetitions [5], [8].

Ultimately, to retain exponential quantum advantage, it might be necessary to conceive applications and algorithms that do not require exponential data loading. Shor's algorithm avoids this problem entirely because of its small input. In quantum chemistry, the Hamiltonian is polynomial in the number of orbitals. Can we find useful applications of quantum linear algebra for "small data" rather than relying on an exponential amount of data?

References

[1] V. Giovannetti, S. Lloyd, and L. Maccone, "Quantum random access memory," *Physical Review Letters*, vol. 100, no. 16, p. 160 501, 2008. arXiv: 0708.1879.

[2] R. Babbush, C. Gidney, D. W. Berry, *et al.*, "Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity," *Physical Review X*, vol. 8, no. 4, p. 041 015, 2018.

[3] G. H. Low, V. Kliuchnikov, and L. Schaeffer, "Trading T-gates for dirty qubits in state preparation and unitary synthesis," 2018. arXiv: 1812.00954.

[4] X. Sun, G. Tian, S. Yang, P. Yuan, and S. Zhang, "Asymptotically Optimal Circuit Depth for Quantum State Preparation and General Unitary Synthesis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2023.

[5] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, "Qubitization of Arbitrary Basis Quantum Chemistry Leveraging Sparsity and Low Rank Factorization," *Quantum*, vol. 3, p. 208, 2019. arXiv: 1902.02134.

[6] J. Welch, D. Greenbaum, S. Mostame, and A. Aspuru-Guzik, "Efficient quantum circuits for diagonal unitaries without ancillas," *New Journal of Physics*, vol. 16, no. 3, p. 033 040, 2014.

[7] V. von Burg, G. H. Low, T. Häner, *et al.*, "Quantum computing enhanced computational catalysis," *Physical Review Research*, vol. 3, no. 3, p. 033 055, 2021. arXiv: 2007.14460.

[8] J. Lee, D. W. Berry, C. Gidney, *et al.*, "Even more efficient quantum computations of chemistry through tensor hypercontraction," *PRX Quantum*, vol. 2, no. 3, p. 030 305, 2021. arXiv: 2011.03494.

[9] S. McArdle, A. Gilyén, and M. Berta, "Quantum state preparation without coherent arithmetic," 2022. arXiv: 2210.14892.

[10] C. Sünderhauf, E. Campbell, and J. Camps, "Block-encoding structured matrices for data input in quantum computing," 2023. arXiv: 2302.10949.

[11] A. V. Ivanov, C. Sünderhauf, N. Holzmann, *et al.*, "Quantum computation for periodic solids in second quantization," *Physical Review Research*, vol. 5, no. 1, p. 013 200, 2023.

[12] D. Camps and R. Van Beeumen, "FABLE: Fast Approximate Quantum Circuits for Block-Encodings," 2022. arXiv: 2205.00081.

[13] G. H. Low, "Halving the cost of quantum multiplexed rotations," 2021. arXiv: 2110.13439.

[14] K. Wan, M. Berta, and E. T. Campbell, "Randomized Quantum Algorithm for Statistical Phase Estimation," *Physical Review Letters*, vol. 129, no. 3, p. 030 503, 2022.

[15] H.-Y. Huang, R. Kueng, and J. Preskill, "Predicting Many Properties of a Quantum System from Very Few Measurements," *Nature Physics*, vol. 16, no. 10, pp. 1050–1057, 2020. arXiv: 2002.08953.

[16] E. Tang, "A quantum-inspired classical algorithm for recommendation systems," in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 2019, pp. 217–228. arXiv: 1807.04271.

[17] S. Aaronson, "Read the fine print," *Nature Physics*, vol. 11, no. 4, pp. 291–293, 4 2015.

Hassle-free Extra Randomness from Quantum State's Identicalness with Untrusted Components

Hamid Tebyanian
Department of Mathematics University of York Heslington, York, YO10 5DD, United Kingdom
hamid.tebyanian@york.ac.uk

***Abstract*—Semi-device-independent (semi-DI) quantum random number generators (QRNGs) are gaining more awareness more and more, presenting a high level of security with an uncomplicated experimental requirement. In this paper, we study a semi-DI protocol based on a minimum error rate of energy-alike coherent states built on the prepare-and-measure scheme with a straightforward experimental requirement where the measurement device is untrusted. Furthermore, the security estimation is based on lower bounding the guessing probability, which is numerically optimized by utilizing semi-definite programming. Finally, a comparison of different encoding and decoding schemes is presented.**

***Index Terms*—Quantum random number generator, Assurance of quantum random number generators**

I. Introduction

Owning high-quality and secure randomness is a necessary step to initiate most of the cybersecurity protocols. The level of security is subject to several theoretical and experimental factors. Accordingly, it is essential to ensure that the random numbers are generated securely to prevent illegal access and obtain testable randomness. In general, random number generators can be classified into three main categories: pseudo-random number generator (PRNG), hardware random number generator (HRNG), and quantum random number generator (QRNG). The PRNG and HRNG are based on deterministic phenomena making them predictable. At the same time, randomness is a fundamental feature of quantum mechanics that originated from its probabilistic nature. Relying on the level of assumptions and experimental conditions, QRNGs can be divided into three sub-groups: device-dependent (DD), semi-device-independent (semi-DI), and device-independent (DI) QRNG. DD QRNGs are easy to implement and performant, while dishonest producers, imperfections, or any deviation from the ideal situation can compromise security. On the other hand, DI QRNGs offer the highest security, but the experimental realization of DI protocol is very challenging, causing them to be less practical. Semi-DI protocols, however, present an excellent trade-off between security and practicality, making them a perfect candidate for practical uses.

II. Protocol

This semi-DI protocol is based on the prepare-and-measure scheme where the preparation device is partially trusted, inspecting the single condition of the protocol, which is transmitting energy-alike states. On the other hand, no requirement is assumed on the measurement device, and it can be treated as a black box.

We consider the simplest encoding technique, binary state preparation, two coherent states represented by two circles in Fig. 1 (A), guaranteeing the states have similar energy; in this case, we can also bind the states' energy to be α-close to the vacuum state, discrimination of such states always comes with an error, see Fig. 1 (B). Basically, the state's indistinguishability detection imposes a minimal rate of unresolved events, namely error probability $P_e \geq 1-\sqrt{1-\delta^2}$, where δ is the scalar product of the states ($\delta = \langle\psi_l|\psi_j\rangle$). As shown in Fig. 1 (B), the error probability is maximum when the states overlap is equal to one ($\mu = 0$), and it decreases with increasing the energy of the states. Otherwise stated, the ambiguity in states discrimination increase when the state's energy decreases; the closer to the vacuum state, the more ambiguity. Note that the states' energy cannot drop to zero as the system becomes single-choice, transmitting vacuum states all the time. The same reasoning applies to more inputs; as long as the state has the energy-alike constraint, the measurement comes with an error.

Suppose that the source takes an input i, chosen independently from the source and the detector, and prepares a physical system in one of the possible quantum states. Later, transmitted to the measurement part, where a detector, which could be either in the continuous or discrete variable domains, e.g., single-photon detector, homodyne, or heterodyne detector, returns an output string o. Randomness can be certified by analysing the input-output probabilities $p\,(i|o)$, given that the states obey the energy-alike constraint. As shown in [1,7,8,9], specific input-output correlations indicate genuine quantum randomness in the sense that the device's output cannot be perfectly predicted, whatever the underlying quantum representation cause it. In spirit, this is comparable to the violation of Bell inequalities which witness genuine randomness independently of the devices' implementation. Owning the measurement outcomes together with the inserted values to the preparation box, we can compute the input-output correlation $p\,(o|i)$:

$$p\,(o|i) = \sum_{\lambda} p\,(\lambda)\left\langle \psi_i^{\lambda} \left|\Pi_o^{\lambda}\right| \psi_i^{\lambda}\right\rangle \tag{1}$$

where ρ_i^{λ} are the propagated states, λ represents the possible strategies of an attacker, and Π_b^{λ} are the POVM determining the measurement method. The conditional min-entropy (CME)

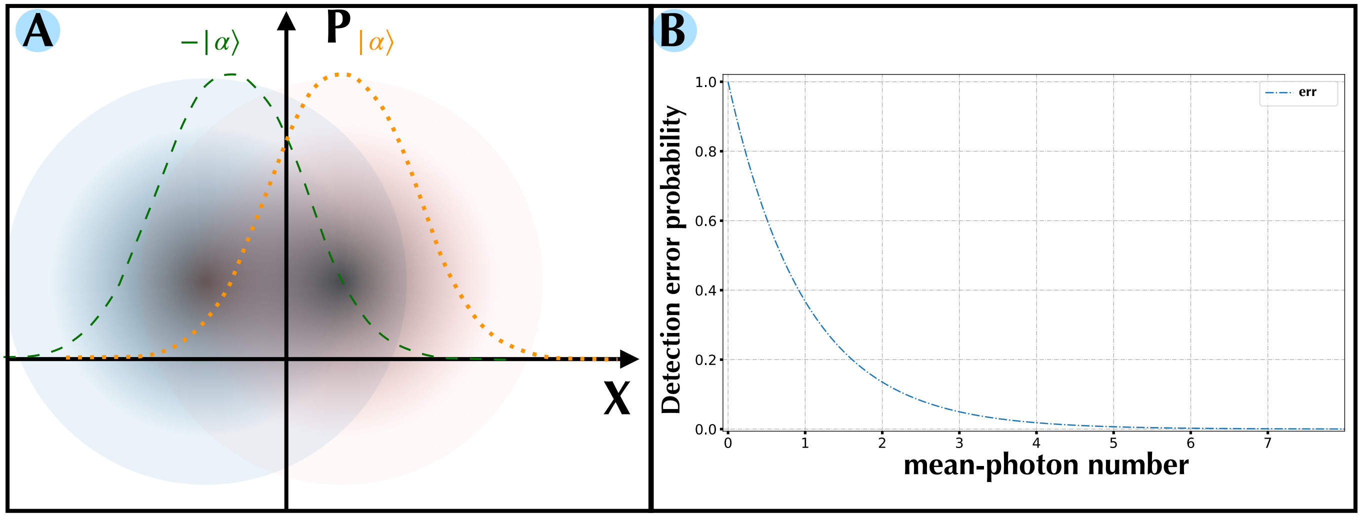

Fig. 1. (A): Schematic representation of coherent state with overlap, the dashed and dotted green and orange lines show the distribution of the states over X. (B): The detection error probability as a function of mean-photon number (μ) of states ($\mu = |\alpha|^2$).

is employed to estimate the system's entropy, put it differently, CME calculates the quantity of extractable genuine randomness which reads:

$$H_{\min} = -\log_2 \left(P_{guess}\right) = \sum_i p_i \sum_\lambda p_\lambda \max\left\{ \sum_o^{m-1} Tr\left[\rho_i^\lambda \Pi_o^\lambda\right]\right\} \tag{2}$$

p_i is the probability of transmitting i, while λ is arbitrary, P_{guess} is the guessing probability, which is the probability that an attacker can guess the outcome, given the input. P_{guess} should be optimized over all possible measurement and preparation strategies, making it complicated to be solved analytically. Following the approach presented in [2,3,4,5,6], we use a numerical tool (semi-definite programming) to solve the optimization problem and estimate the amount of extractable randomness.

III. Preparation and Measurement

The preparation device transmits quantum states with limited energy; this constraint can be seen in the context of energy-bound or overlap assumptions depending on the user's choice. The energy bound is a tighter bound which imposes an inevitable overlap of the prepared states. In any case, the source can be a weak coherent source or a time-bin single-photon source as long as the experiment energy-alike condition is satisfied.

Here we demonstrate the phase encoding scheme exploiting a weak coherent light, meaning that the states are encoded based on their phase. This provides the possibility to increase the number of inputs straightforwardly, as shown in Fig. 2 (A); infinite possible states can be encoded in this way, while the states are located within limits, dashed circle. On the measurement side, in general, any scheme can be employed; here, we study heterodyne detection as it gives information on both light field quadratures simultaneously, enabling tracking the states' phase. The heterodyne detection describes the probability density of getting detection proportional to $\frac{(X+iP)}{2}$ from an optimal simultaneous measurement of field quadratures X and P. This kind of measurement is undoubtedly non-ideal, considering that the field quadratures X and P do not commute. The heterodyne detection corresponding POVM reads;

$$\Pi\left(x_\phi\right) = |\beta\rangle\langle\beta| \tag{3}$$

Where $|\beta\rangle$ is the coherent state with complex amplitude β. Having the POVM, we can compute the conditional probabilities:

$$p(\alpha) = 1/\pi \int ||\langle\alpha|\beta\rangle|^2 \, d\beta^2 \tag{4}$$

Fig. 2(B) represents the conditional min-entropy as a function of the number of outcomes for the heterodyne detector for binary and ternary encoding schemes. As shown, the entropy improvement for the binary case, even for higher outcomes, is negligible, but for the ternary case, the gain is more evident.

Note that increasing the number of outcomes can be done in the post-processing stage without touching the actual experimental setup. Indicating that more randomness can be extracted from the same optical device only by changing the data processing stage.

IV. Conclusion

In conclusion, this paper studies a semi-DI QRNG with various encoding and decoding schemes, particularly phase encoding and heterodyne detection schemes were investigated in detail. It is shown that by increasing the number of inputs and outcomes, the extractable randomness increases accordingly, meaning more accessible randomness needles for changing any experimental component.

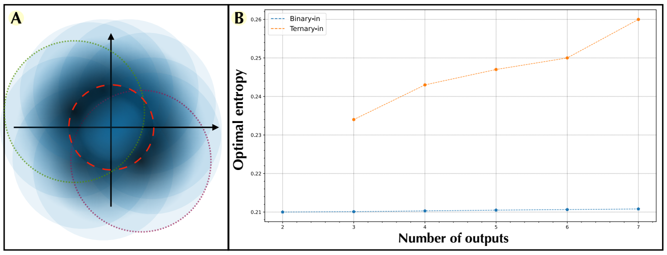

Fig. 2. (A): Encoding technique using weak coherent phase; in this way, one can generate infinite inputs using different phases to encode the states. The central dashed circle represents the experiment constraint, and the two dotted lines exemplify the binary phase keying scheme. (B): The conditional min-entropy as a function of the number of outcomes for heterodyne detection with binary and ternary phase shift keying encoding schemes.

V. Acknowledgement*

This work is supported by the Innovate UK Industrial Strategy Challenge Fund (ISCF), project 106374-49229 AQuRand (Assurance of Quantum Random Number Generators).

References

[1] Rusca, D., Tebyanian, H., Martin, A., Zbinden, Fast self-testing quantum random number generator based on homodyne detection, Applied Physics Letters 116 (2020).

[2] Avesani, M., Tebyanian, H., Villoresi, P., Vallone, G. Unbounded randomness from uncharacterized sources. *Communications Physics*. **5**, 273 (2022).

[3] Tebyanian, H., Zahidy, M., Avesani, M., Stanco, A., Villoresi, P. & Vallone, G. Semi-device independent randomness generation based on quantum state's indistinguishability. *Quantum Science And Technology*. **6**, 045026 (2021,9), https://doi.org/10.1088/2058-9565/ac2047

[4] Tebyanian, H., Avesani, M., Vallone, G., Villoresi, P. Semi-device-independent randomness from d-outcome continuous-variable detection. *Phys. Rev. A*. **104**, 062424 (2021).

[5] Brask, J., Martin, A., Esposito, W., Houlmann, R., Bowles, J., Zbinden, H. & Brunner, N. Megahertz-Rate Semi-Device-Independent Quantum Random Number Generators Based on Unambiguous State Discrimination. *Phys. Rev. Applied*. **7**, 054018 (2017,5), https://link.aps.org/doi/10.1103/PhysRevApplied.7.054018

[6] Tebyanian, et al., Time-bin Quantum Random Number Generator with Uncharacterized Devices.

[7] Tebyanian, H., Avesani, M., Vallone, G. & Villoresi, P. Semi-device-independent randomness from d-outcome continuous-variable detection. *Phys. Rev. A*. **104**, 062424 (2021,12), https://link.aps.org/doi/10.1103/PhysRevA.104.062424

[8] Avesani, M., Tebyanian, H., Villoresi, P. & Vallone, G. Semi-Device-Independent Heterodyne-Based Quantum Random-Number Generator. *Phys. Rev. Applied*. **15**, 034034 (2021,3), https://link.aps.org/doi/10.1103/PhysRevApplied.15.034034

[9] Zahidy, M., Tebyanian, H., Cozzolino, D., Liu, Y., Ding, Y., Morioka, T., Oxenløwe, L. & Bacco, D. Quantum randomness generation via orbital angular momentum modes crosstalk in a ring-core fiber. *AVS Quantum Science*. **4**, 011402 (2022),

Approximate Computing for Noise Resilient Resource Optimized Quantum Circuits

Himanshu Thapliyal*, Joseph Clark, and Bhaskar Gaur
University of Tennessee, Knoxville, TN, USA; *Corresponding Author: hthapliyal@utk.edu

Topic: Compilation; Error Correction and Mitigation;

1 Challenge

Fully fault-tolerant quantum computation will take a significant amount of resources. Therefore, one of the current focuses in quantum computation is to establish the utility of small-scale, error-prone, or "noisy intermediate-scale quantum" (NISQ), machines. In NISQ machines, the application of quantum gates, as well as measurement operations, can introduce errors. Further, the amount of time a qubit register can maintain its state (coherence time) is short. So, to enable progress in quantum computation, there are two challenges: (i) develop strategies at various levels of abstraction to increase the noise resilience of the quantum circuits in NISQ Machines, (ii) develop the resource-efficient design of quantum circuits that can have applications from NISQ to Fault-tolerant quantum computation. With the ultimate goal of enabling practical applications of quantum computing, there is a need to develop new circuit-level techniques and design automation tools for mitigating errors and optimizing the use of computational resources in quantum circuits.

2 Opportunity

Approximate computing is a novel computing paradigm that produces imprecise results by relaxing the need for fully precise or completely deterministic operations; such error-tolerant applications include multimedia, data mining, and image processing [1]. Approximate computing often utilizes statistical properties of data and algorithms to trade away quality for improvements in other figures of merit (such as energy, area and power reduction in classical computing). In quantum computing, approximate computing can trade circuit accuracy to optimize other parameters such as noise fidelity, quantum depth, and resource utilization. This approach is particularly useful in applications domains of quantum computing where exact results are not required, such as machine learning, signal processing, and image processing. Among quantum circuits, quantum arithmetic circuits serve as the building blocks for complex applications. Creating their approximate libraries can be particularly helpful in error-tolerant applications such as quantum image processing and quantum machine learning. Also, there is a need for design automation tools that can produce: (i) approximate quantum circuits that achieve higher accuracy than the baseline circuit when executed on NISQ machines, (ii) approximate quantum circuits that are resource efficient compared to the baseline circuit when executing on Fault-tolerant quantum machines. Table I provide details of the characteristics and implementation of several approximate circuit design techniques. It contains the name of the method, a reference to the work the technique was introduced in, and a description of the technique.

Nature-inspired algorithms, many of which are inherently approximate, are already very common in quantum synthesis, due to their adaptability and tendency to find good solutions to complex problems. These algorithms have also been applied to great effect in approximate synthesis applications [2], [3]. Existing works in this direction are still only able to handle relatively small problems, and the scalability of these algorithms is generally not well-characterized (although it is generally expected to be exponential in the number of qubits). Further, error-aware design takes into account the error rate of underlying hardware when designing and laying out a quantum circuit. This is a useful design methodology for NISQ design in general but is particularly beneficial for noise-resilient approximate design. Unfortunately, this methodology is still in its infancy, but existing work has shown very promising results [4].

Further, the development of standardized quantum libraries for approximate math functions will help domain science experts to experiment and develop new noise-resilient and resource optimized quantum simulation algorithms and applications. Recently, five designs of approximate quantum adders designed to reduce depth while making them noise-resilient at the same time are proposed. For various noise models compared to exact designs of quantum ripple carry adder, the approximate quantum adders have improved fidelity ranging from 8.34% to 219.22% [5]. Creating these approximate math libraries of all arithmetic and scientific functions can be particularly helpful in error-tolerant applications such as quantum image processing and quantum machine learning.

There is an opportunity to advance quantum approximate circuit design by developing new design methods and tools able to generate good results for large circuits. As an example, the Tree-based Directed Acyclic Graph (TDAG) [6] partitioning for quantum circuits, a novel quantum circuit partitioning method that partitions circuits by viewing them as a series of binary trees and selecting the tree containing the most gates is a step in this direction. Novel nature-inspired algorithms should be invented and applied to address the increasing size of quantum circuits and to succeed in producing

TABLE 1: Quantum Approximate Circuit Design Techniques

Name/Method	Reference	Approximation Technique
QUEST	[2]	Partition circuit into blocks, perform approximate synthesis on each block, combine approximate blocks using simulated annealing to form many dissimilar approximations
QAQC	[3]	Determine circuit structure using simulated annealing, find gate parameters using a gradient-based or one of two gradient-free methods with a Hilbert-Schmidt test; for large circuits, use an ansatz to determine circuit structure to avoid the large structure search space, use a localized variant of the Hilbert-Schmidt test which is more stable for large circuits
Optimization via Energy Minimization	[7]	An ansatz forms the structure of the circuit, and a constant input value has been provided; use energy minimization to determine the parameter values for the ansatz, use the "lure" method to iteratively produce a series of mappings which satisfy the requirements
MPS Approximation	[8]	Given a target matrix product state and a fixed circuit structure, set the parameters in the circuit using gradient descent with a localized error function
NACL	[4]	Given a cost function based on the circuit specification and a noisy simulation of the target hardware, produce an approximate circuit which minimizes the cost function, including the effect of hardware noise

circuit structures where other methods fail. Further, the integration of error awareness into approximate circuit design has the potential to improve the circuit performance on NISQ machines by directly focusing on result fidelity.

3 Assessment

The following metrics can be used to assess the reliability of the proposed quantum approximate circuits: (1) ***Fidelity:*** The fidelity computes the closeness of the two quantum states. Fidelity equal to 1 means that two states are equal. An increase in noise reduces the fidelity because there is a deviation in the actual output. Therefore, the higher the fidelity of a quantum circuit, it is better in terms of reliability. (2) ***Probability of Successful trial (PST):*** PST evaluates the reliability of NISQ applications by computing the ratio of the number of error-free trials to the total number of trials. PST metric is widely used as evident from the existing works. (3) ***Inference Strength (IST):*** IST evaluates the ratio of the frequency of error-free output to the number of most frequently occurring erroneous output. If IST > 1, the system will infer the output as correct otherwise the incorrect answer(s) will dominate the correct answer. Further, we can measure the impact of introducing approximation on the accuracy of the quantum application by using the error metric of **Normalized Mean Error Deviation (NMED)**. NMED is calculated by taking the mean of the absolute error deviation divided by the theoretical maximum value for the given input set. A figure of merit called **Quantum State Fidelity Ratio (QSFR)** [9] to measure the closeness of the correct output to the top output is also proposed recently. QSFR quantifies the degree of improvement needed to advance the intended output to the top position. In addition, circuit approximations can be evaluated using **Total Variation Distance (TVD)** and the **Jensen-Shannon Divergence (JSD)**.

Resource cost measures that can be used in evaluating quantum approximate circuits include but are not limited to the number of qubits, the number of gates, the depth that is the maximum number of gate layers in a quantum circuit, etc.

4 Timeliness

In the current stage of quantum computing, we have quantum hardware that scales better but with persistent noise issues. More qubits open the field of quantum computing to more applications. As more applications utilize quantum computers, noise fidelity gains priority. Approximate computing can be a promising way forward to increase noise fidelity, as it reduces depth together with complex dependencies among qubits. Resulting in quantum approximate circuits showing greater flexibility to physical layout. As approximate computing inherently introduces some error, it is suitable to enable inherently error-tolerant quantum applications. Quantum approximate computing is uniquely useful for NISQ applications since it improves accuracy by reducing resource usage and improving speed. When fault-tolerant quantum computers are ready to be deployed, the quantum approximate design will remain useful as they can be used to save quantum resources and improve depth (timing/delay).

References

[1] W. Liu and F. Lombardi, *Approximate Computing*. Springer, 2022.

[2] T. Patel, E. Younis, C. Iancu, W. de Jong, and D. Tiwari, "QUEST: systematically approximating quantum circuits for higher output fidelity," in *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2022, pp. 514–528.

[3] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, "Quantum-assisted quantum compiling," *Quantum*, vol. 3, p. 140, 2019.

[4] L. Cincio, K. Rudinger, M. Sarovar, and P. J. Coles, "Machine learning of noise-resilient quantum circuits," *PRX Quantum*, vol. 2, no. 1, p. 010324, 2021.

[5] B. Gaur, T. Humble, and H. Thapliyal, "Noise-resilient and reduced depth approximate adders for nisq quantum computing," in *ACM GLSVLSI 2023*.

[6] J. Clark, T. Humble and H. Thapliyal, "Tdag: Tree-based directed acyclic graph partitioning for quantum circuits," in *ACM GLSVLSI 2023*.

[7] T. Jones and S. C. Benjamin, "Robust quantum compilation and circuit optimisation via energy minimisation," *Quantum*, vol. 6, p. 628, 2022.

[8] M. B. Dov, D. Shnaiderov, A. Makmal, and E. G. D. Torre, "Approximate encoding of quantum states using shallow circuits," *arXiv preprint arXiv:2207.00028*, 2022.

[9] B. Gaur, E. Munoz-Coreas, and H. Thapliyal, "A logarithmic depth quantum carry-lookahead modulo (2n - 1) adder," in *ACM GLSVLSI 2023*.

Title: Quantum Privacy Preserving Artificial Intelligence (QPPAI)

Authors: **Huan-Hsin Tseng** (BNL, htseng@bnl.gov), Katheleen Hamilton (ORNL, hamiltonke@ornl.gov), Meng Yue (BNL, yuemeng@bnl.gov), Samuel Yen-chi Chen (ycchen1989@gmail.com), Shinjae Yoo (BNL, sjyoo@bnl.gov)

Topic: Algorithm and Models

Challenge: Quantum Artificial Intelligence (QAI) is a rapidly growing area with both theoretical and empirical quantum advantages. However, QAI has not been considered the privacy aspect of QAI much. Here are use cases for secure Quantum AI:

- **Biomedical Applications** Modern AI/ML techniques provide a suitable toolset for analyzing large volumes of medical data including medical images and electronic health records (EHR). The security of these sensitive data is of the highest priority. Quantum enhanced AI/ML techniques can be applied to institutes such as Veterans Affairs (VA) and the National Institute of Health (NIH) where a significant amount of data is stored and yet to be analyzed.
- **Smart Grid** The rapidly increasing intermittent, undispatchable renewables and electrification of transportation, building, etc. poses a great threat to the reliability, security, and stability of the power grid. Situational awareness becomes increasingly critical to power operation and control, calling for innovative data-driven and AI applications such as state estimation, prediction of grid dynamic trajectories and stability, and data-driven modeling and control strategies while preserving the privacy of data of different utility owners. QPPAI can provide a promising solution to such applications.
- **User Facilities** Some user facilities such as National Synchrotron Light Source (NSLS-II) and Center for Functional Nanomaterials (CFN) provide significant help in designing novel materials. The process of a large volume of high-resolution images poses great challenges to classical computing. Quantum computing is a promising technology to provide high-throughput analytics for these valuable data. However, certain users rely on these facilities to develop advanced and proprietary materials, which should be protected as trade secrets and not disclosed to unauthorized third-party.
- **Nuclear Safeguard** The nuclear facilities in each country need to be inspected regularly and nuclear safeguard data is important to build automated safeguard application but each country has different policy and data sharing are often not recommended. Pooling such safeguard data together could help us to build more powerful and secure nuclear safeguard AI models using quantum computing.

One of the common features of these successful ML models is that they are data-driven, from biomedical image analysis to nuclear safety AI models and high-resolution images in material science. To build a successful machine learning model a huge amount of data is required. Although there are several public datasets for research purposes, most advanced and personalized models largely depend on the collected data from the edge and IoT devices, such as mobile phones and other personal/commercial data (*e.g.*, medical records, browsing habits, images from proprietary synthesized materials, nuclear power plant operations and etc). For example, ML/DL also succeeds in the field of medical imaging, speech recognition, autonomous vehicles, etc. These fields rely critically on the massive dataset collected from the population and these data should not be accessed by unauthorized third-party. With the rapidly growing field of quantum AI/ML technologies, we expect these emerging techniques to be employed in various biomedical, scientific, commercial, and nuclear security scenarios which largely depend on private or sensitive datasets. There are three major challenges to guiding secure quantum AI:

- **Ensuring no physical data sharing while training with all available data.** Modern machine learning models heavily rely on massive datasets, which inevitably contain personal or sensitive information. For utilization of QAI with existing cloud-based quantum computers, it is desirable to train AI models without sharing users' personal or sensitive data in quantum computers.
- **Secure data communication and storage.** The widely used channel to exchange data with the cloud service could be compromised, leading to the leakage of high-value personal or commercial data. Even if a communication channel can be secured, cloud service providers contain potential risks as malicious adversaries can potentially infiltrate computing infrastructures.
- **Ensuring privacy from the trained model.** Keeping private data is not enough. The adversaries can still deduce a single data entry by attacking the trained model. It is therefore urgent to build

models immune to such attacks. One of the difficulties of such privacy-preserving optimization is that it usually affects the model performance. It is still unknown to what extent such optimization procedures would affect QML models and whether there are actually quantum advantages.

Opportunity: We propose mainly three research directions to cope with the aforementioned quantum secure computing challenges:

- **Quantum Federated Learning (QFL)**: QFL [1] is the computational framework consisting of an array of quantum computers connected together to solve the problem of interest. The main feature of QFL is that the data is kept in one quantum computer but not shared with others. The potential research direction is to study an efficient way to distribute the computational load among different quantum computing machines, each with heterogeneous computing capabilities and data repositories. In addition, the method for efficient model aggregation with both *classical and quantum network* are of high interest.
- **Secure Quantum Multi-Party Computation**: Homomorphic Encryption (HE) allows computation under encrypted data while the results of computation remain encrypted [2]. Secure Quantum Multi-party computation provides secure protocols for the communication and computing between multiple quantum computers. In particular, how to securely communicate data between multiple and decentralized classical and quantum computers are of high interest as in the near-term, the hybrid quantum-classical paradigm will play a major role in the quantum computing community. This concept is to be integrated with our proposed quantum federated learning.
- **Quantum Differential Privacy (QDP)**: Differential Privacy (DP) [3] ensures each individual data leakage from the trained model and we propose the quantum version of DP for AI models, so that the trained model would not be attacked. Promising research directions include the efficient optimization algorithms for such quantum-DP models as well as the aggregation methods for distributed quantum-DP training. In general, DP training will degrade model performance and therefore it is an interesting research direction to investigate whether the quantum advantage suppresses DP training loss or not.

We expect the combination of these three techniques would greatly boost the security, scalability and reliability of the QML application on NISQ machines and beyond.

Assessment: We identify the following metrics to quantify success and evaluate potential solutions.

- What are the quantum advantages and overheads of quantum federated learning over classical federated learning?
- Can a quantum differential privacy trained AI model be equipped with a better probabilistic privacy guarantee than the classical one with the same level of accuracy?
- Can data be encoded into quantum subsystems in a way to avoid privacy risks (*e.g.*, using Machine Unlearning [4])?

Probing the potential of quantum advantages over classical PPAI is essential. Under general application, the improvement of security protocol and resilience during the privacy attack is a key assessment criterion.

Timeliness or maturity: With the advancement of quantum internet, quantum sensing technologies, and rapidly growing logical qubits, the quantum research community is advised to commence on QPPAI investigation as the privacy is one of the stumbling blocks of the actual deployment of QAI applications.

References

[1] S. Y.-C. Chen and S. Yoo, “Federated quantum machine learning,” *Entropy*, vol. 23, no. 4, p. 460, 2021.

[2] Y.-J. Zhang, J.-W. Liu, T. Shang, and W. Wu, “Quantum homomorphic encryption based on quantum obfuscation,” in *IWCMC*, pp. 2010–2015, IEEE, 2020.

[3] W. M. Watkins, S. Y.-C. Chen, and S. Yoo, “Quantum machine learning with differential privacy,” *Scientific Reports*, vol. 13, no. 1, p. 2453, 2023.

[4] M. Chen, Z. Zhang, T. Wang, M. Backes, M. Humbert, and Y. Zhang, “When machine unlearning jeopardizes privacy,” in *Proceedings of the 2021 ACM SIGSAC*, CCS ’21, p. 896–911, 2021.

Resources for Heterogeneous Quantum Networking: Quantum Frequency Conversion

Raju Valivarthi[*], Prathwiraj Umesh, Cristián Peña, Si Xie, Neil Sinclair, Panagiotis Spentzouris, Prem Kumar, Maria Spiropulu

May 2, 2023

Topic: *co-design and integration*

1 Introduction

Quantum networks are at the forefront of scientific research with promising applications in secure communications, advanced computing, and sensing. One of the major challenges of co-designing a scalable and multi-purpose quantum network is to reversibly map quantum information between stationary qubits (typically matter) and flying qubits (those that traverse fiber or free-space channels), or between flying qubits (when using photons with different encoding schemes). Furthermore, it also requires the mapping of qubits between photons of different energies, a closely related task. Moreover, photons must be indistinguishable in their degrees of freedom to interfere, which is the necessary requirement to create remote entanglement. Quantum frequency conversion (QFC) is an essential resource to solve this problem by enabling the conversion of photons between different frequencies while, if performed in an efficient and noise-less fashion, preserving the quantum state of the original photons. We argue that QFC is a key technology to enable communication between multiple qubit platforms, effectively enabling heterogeneous distributed quantum computing and sensing.

A typical approach for QFC [1] uses parametric wave mixing based on $\chi^{(2)}$ and $\chi^{(3)}$ optical nonlinearities in solid-state crystals and fibers. This approach is compatible with a variety of wavelengths, allows broadband conversion, operates at room temperature, and does not suffer from strong noise processes introduced by near-resonant transitions, cf. four-wave mixing in atomic ensembles [4]. This approach requires combining a strong classical pump with the input photon both to bridge the frequency gap and permit efficient conversion. However, the strong optical pump may generate undesired noise photons due to far off-resonant interactions and photon emission due to additional nonlinear interactions with the pump (see Challenge below). To mitigate the noise, an alternative, yet closely related, approach to perform the frequency conversion involves creating entangled photons with the pump and performing quantum teleportation of the photonic qubits. However, this process is somewhat inefficient due to low probability of entanglement creation. These trade-offs must be considered, and optimized for, to operate the quantum network for a particular application (see Opportunity below).

2 Challenge

Although there has been significant efforts in the research community to develop QFC technologies, an efficient, low-noise QFC interface has not been developed. Efficient conversion requires careful control of the impact of material dispersion on the relevant photon fields, and thus QFC has been explored in a variety of materials including KTP, BBO, and $LiNbO_3$. Cavity and waveguide based approaches

[*]Corresponding author: rajuval@caltech.edu

have been considered to enhance the QFC conversion efficiency, as well as the possibility to engineer the phase matching using polarization matching and periodic poling. Further challenges involve efficiently interfacing the converter with the channel and matter qubits. Yet, the ultimate goal of realizing efficient conversion of a photon into a particular spatio-temporal mode without additional noise is still outstanding. Until now, efforts have focused on optimizing over an ensemble of certain parameters [2] like internal conversion efficiency, fibre coupling to the waveguides, mode properties of frequency-converted photons, and noise rates. However, research needs to be pursued towards optimization of figures of merit that enable an heterogeneous quantum network.

3 Opportunity and Assessment

A survey of the research community and relevant industry partners must be performed to identify the most relevant approaches for QFC towards the development of quantum networks. In particular the QFC approach must be optimized for the overall functioning of the network, which involves taking into account specific protocol designs and functionalities, flying and stationary qubit properties including hardware specifics, detectors, and use-cases e.g. direct vs teleportation-based conversion. Realization of the needed QFC will require extensive technical engineering advancements such as efficient device design, high-quality crystal or fiber growth, high efficiency fiber coupling and miniaturization of devices [3]. This requires collaborations from all the key research and relevant industry partners. Successful collaboration will be able to deliver theory-guided QFC hardware for each of the targeted conversions.

The assessment of quantum frequency conversion (QFC) devices involves evaluating several critical parameters, including conversion efficiency, noise introduced, compactness, ease of integration, and resource requirements. High conversion efficiency is desirable to minimize the loss of photons during conversion, and noise must be minimized to maintain the fidelity of quantum communication protocols. Compactness and ease of integration are important to reduce the size of the overall system and enable integration with other quantum technologies. Resource requirements such as power consumption, cooling, and maintenance should also be considered to ensure practicality and scalability. By evaluating these parameters, researchers can identify promising QFC devices for further development and advancement towards the widespread adoption of QFC technologies.

4 Timeliness or Maturity

The timeliness of QFC research and development is significant, as it is an essential resource needed for the co-design of a multi-purpose quantum network. As quantum networks continue to gain attention and funding, there is an increasing demand for QFC devices that can efficiently convert photons between different frequencies while preserving their quantum state. Continued research and collaboration between academia and industry will be essential for realizing this potential and bringing QFC devices to market.

References

[1] Kumar, P., 1990. Quantum frequency conversion. Optics letters, 15(24), pp.1476-1478.

[2] Fejer, M.M., 1994. Nonlinear optical frequency conversion. Physics today, 47(5), pp.25-33.

[3] Moody, G., Sorger, V.J., Blumenthal, D.J., Juodawlkis, P.W., Loh, W., Sorace-Agaskar, C., Jones, A.E., Balram, K.C., Matthews, J.C., Laing, A. and Davanco, M., 2022. 2022 Roadmap on integrated quantum photonics. Journal of Physics: Photonics, 4(1), p.012501.

[4] Becerra, F.E., Willis, R.T., Rolston, S.L. and Orozco, L.A., 2008. Nondegenerate four-wave mixing in rubidium vapor: The diamond configuration. Physical review A, 78(1), p.013834.

Short-Term Quantum Network-Assisted Applications

Dinesh Verma (IBM, dverma@us.ibm.com), Mark Ritter (IBM, mritter@us.ibm.com), Shiqiang Wang (IBM, wangshiq@us.ibm.com), Gabriella Carini (BNL, carini@bnl.gov), Dimitrios Katramatos, Brookhaven National Laboratory, dkat@bnl.gov, Bogdan Neculaes (GE Research, Bogdan.Neculaes@ge.com), Stephen Bush, (GE Research, Stephen.Bush@ge.com), Paul Dabbar, (Bohr Quantum, paul.dabbar@bohrquantum.com), Eden Figueroa, SUNY Stony Brook, eden.figueroa@stonybrook.edu.

Topic: Quantum Networks Applications & Quantum Internet Stack

1. Challenge: Define classical applications than can benefit from quantum communication.

Quantum Networking and Computing are perceived as futuristic technologies with commercial adoption and use-cases expected beyond the 5-year horizon. While that holds true for several topics that fall within the scope of quantum computing and networking, there are many applications of quantum networks that could be attainable in the short term. These short term applications can be developed using classical computers assisted by quantum network connections. Within this context, quantum communication attributes can enable new classes of applications, as well as simplify the operation and logistics of existing applications. Our hypothesis is that several such quantum-network-assisted applications exist, and a concerted effort is necessary to further identify, develop, and demonstrate such applications.

Among the unique attributes of quantum networks, that quantum network-assisted applications can benefit from, we list: (i) *Unbreakable Security*, which is based on quantum principles such as superposition, entanglement, and non-cloning, enabling the generation of secure keys, usable in one-time pad or symmetric encryption, (ii) *True randomness*, which is used in many quantum communication protocols to generate quantum-certified streams of random numbers and qubits and provides the possibility to be used in distributed Monte Carlo simulations, and (iii) *Teleportation-Assisted Wide Area Connectivity*, which is established by distributing multipartite entangled states across many nodes, and can be used to drive teleportation events servicing links between classical computers.

To benefit from these quantum attributes in conjunction with current communication networks, a few key challenges must be overcome, such as: (1) *Slow Speed of Quantum*: current quantum-protected key generation and entanglement distribution schemes can achieve rates of a few MHz, which must be matched to operate together with classical services that operate at GHz speeds, (2) *Denial of Service for Quantum*: quantum data flow can be disrupted by a variety of attacks, including blinding and noise addition, which must be mitigated to be part of larger and more reliable classical communication services, and (3) *Infrastructure Requirements for Quantum*: quantum communication hardware is restricted in the short-term to sites having large laboratory facilities, thus, it is imperative to explore the deployment of quantum equipment in communication data centers, to facilitate coexistence and parallel operation with classical optical networks.

2. Opportunity: Using the the socket layer abstractions together with quantum attributes.

We believe that these challenges must be addressed by bringing together a cross-disciplinary set of researchers. To date, the focus of quantum networks has been on the development of the physical layer. While that effort should continue and has made substantial progress, building a suite of practical applications that exist on the upper layers of the quantum network protocol stack will require additional efforts from networking software engineers, as well as computer science theorists and programming language experts. Along these lines, there exists a need to assess and define new network abstractions that by design include quantum communication attributes as a part of the operation of classical networks. These abstractions can be envisioned in three different configurations: (i) a system interconnecting multiple classical computers using in-network quantum attributes, (ii) a quantum internet network interconnecting multiple quantum computers, and (iii) a hybrid quantum/classical network interconnecting a mixture of classical and quantum computers. Among the applications that can be built from these abstractions, we mention distributed quantum computing and blind quantum computing as known perspective applications of (ii) and (iii). However, due to the lack of scalable operational quantum computers, these applications remain on the long-term horizon. Nonetheless, we see a clear opportunity to develop (i), by using the socket layer model, including the verifiable byte stream, unreliable datagram, and secure socket abstractions, in combination with the quantum network attributes described above.

Quantum Network Assisted Applications

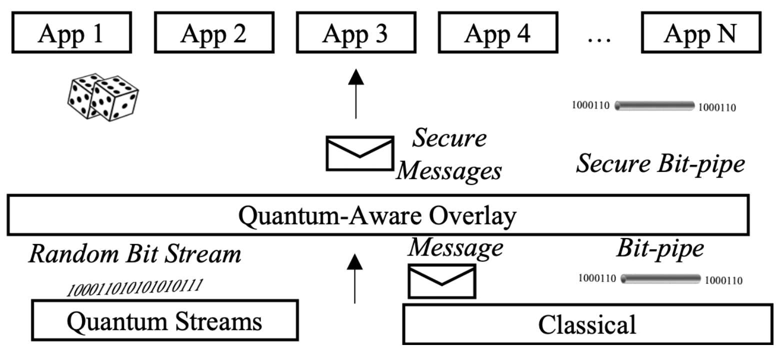

In this configuration, a possible set of abstractions is shown in the diagram on the right. Here, quantum channels and classical channels work in parallel, supporting the communication of quantum streams and classical communication packets. We envision a quantum-aware overlay, created as an abstraction on top of these channels, which can combine quantum-secure keys and PQC techniques to offer secure communications or just take advantage of the true randomness of the quantum streams to provide truly random numbers to applications on the network [1-2]. Quantum-assisted software applications can then be written atop this abstraction.

3. Assessment: Development and demonstration of quantum-network-assisted applications.

Success would entail successful development and demonstrations of quantum-network-assisted applications running on hybrid classical/quantum networks. Among the applications that can benefit from the abstractions defined above, we envision enterprise applications running atop a messaging protocol or an Enterprise Service Bus, large scale Monte Carlo simulations, parallelized Las Vegas algorithms, probabilistic machine learning approaches, and genetic algorithms, as well as enhanced security mechanisms such as moving target defense, and enhanced optimization methods. Additionally, this abstraction can also be useful to implement more advanced applications, such as quantum secure direct communication [3] and secure optical communication using quantum alarms [4]. These applications can find first use-cases and short-term practical use in ensuring security of several energy delivery sites that need to provide a public facing presence. These sites may need to exchange confidential information with each other, while also avoiding attacks from malicious insider attacks from within or externally.

4. Timeliness/maturity: evaluation of improvement/advantage with respect to classical networks.

Qubits streams and entanglement distribution are at an early stage of development for long-distance operation. However, first generation quantum networking systems are ready to begin integration with software systems running over classical computers. Maturity of such novel systems must be evaluated using empirical or theoretical analysis of why the approach assisted by quantum networks is better than the approach using only classical networks. The improvement/advantage may be measured in terms of metrics such as application latency, application throughput, or application ease of deployment.

References:

[1] Software Architecture for Operation and Use of Quantum Communications Networks, submitted to. IEEE Services Quantum Software 2023.

[2] Configuration Distribution using Quantum Networks, submitted to IEEE Quantum Week, 2023.

[3] Quantum Secure Direct Communication, Light: Science & Applications 10, 183 (2021).

[4] Secure optical communication using a quantum alarm, Light: Science & Applications 9, 170 (2020).

Classical Simulations of Quantum Computers

Yan Wang[1,†], Joseph Wang[1], Phil Lotshaw[1], Ryan Bennink[1]
[1]Oak Ridge National Laboratory, Oak Ridge, TN. [†]wangy2@ornl.gov

Topic: algorithms—classical algorithms and software systems to simulate quantum computers and networks, including tensor network and Monte Carlo simulations.

Challenge: As remarked in the pre-workshop report[1], quantum computers (QCs) are expected to violate the strong Church–Turing thesis[2] and hence have the potential to outperform any classical computers by using exponentially fewer resources. In general, the exponential cost makes simulating QCs a formidable task with today's classical computers. For example, using 1 byte (B) (so-called 8-bit mini float) for each matrix element of a 40-qubit quantum state vector, the memory cost alone would be comparable to 2^{40}B∼1TB, the typical size of data used to train large language models (LLMs). In addition to this exponential wall, the larger the system is, the higher precision the matrix elements need to be represented in computation to obtain the correct final answer—resulting in the insurmountable computing cost referred as Van Vleck catastrophy[3]. Nevertheless, if the problem instances and quantum processes executed on QCs have certain *structures* that can be exploited, such as constrained entanglement growth and/or conserved symmetries, they are simulatable with realistically available classical resources measured in terms of the length of compute time, the amount of memory, and the number of processor units (CPUs, GPUs, TPUs, etc.). How to make the best use of the classical resources to simulate these *structured problems* on QCs at a nontrivial scale as large as possible, remains a great challenge. Conquering this Herculean task[4] is invaluable to the design, characterization, and verification of quantum hardware and to empirical prototyping and evaluation of quantum algorithms and protocols implemented on the hardware.

Opportunity: To illustrate the exploitation of problem structure for efficient algorithm and computation, consider the elementary problem of summing the arithmetic progression $\sum_{k=1}^{n} k = n(1+n)/2$. Directly adding each term one by one has $O(n)$ complexity while the sum formula on the right has $O(1)$ complexity. The structure of this problem is the symmetry of the series: when all terms are plotted as dots along the real axis, they form a lattice with inversion symmetry with respect to the center, a property that can be used to derive the $O(1)$-complexity sum formula. We believe that exploring and exploiting the problem structure is the guiding principle of opening new opportunities for simulating QCs. For the problem of simulating quantum circuits executed on large QCs ($\gtrsim 100$ qubits), currently, the only viable methods are matrix product state (MPS) simulator based on tensor network (TN) and stabilizer simulator: the former exploits the confined entanglement in space and the latter stabilizer group symmetry. On the other hand, simulating quantum dynamics described by general quantum circuits with minimal bias (less approximation) *and* manageable resource is still a hard problem for these simulators or any others. Finally, simulations of noisy quantum circuits are especially expensive, as they require simulating or sampling from a density matrix (a difficulty also appears in simulating quantum dynamics), with resource costs that significantly exceed similar simulations of noiseless circuits. However, to assess progress, evaluate performance, and understand behaviors of modern quantum devices and algorithms with tens and hundreds of qubits, we must find ways to model and simulate large, noisy quantum circuits.

Assessment: A performant classical algorithm to simulate quantum dynamics should scale well in space (memory cost) and time complexities. Even though there have been efforts to mitigate the memory issues by parallel programming with smart memory allocation in a high-performance environment, better algorithm development is still needed to remove problems such as proper basis truncation in the history of time evolution to avoid keeping track of wave amplitudes at all times, which is the current bottleneck for the approach. In addition, handling mixed states in the simulation is necessary for including thermal and environmental decoherence effects. A possibility is to adapt the quantum trajectory approach[5,6] that has been proven equivalent to the density-matrix approach and can incorporate fast quantum noise and the collapse of the quantum state into other states to different environments through measurement.

Although tensor-network simulations have been used for example in laptop simulations of random circuits like the quantum supremacy experiments, with comparable fidelity to those experiments[7], the fidelity metric that is used to truncate these tensor networks does not have a straightforward relationship to computationally expensive noise models that are typically used to characterize devices and algorithm performance. Thus, tensor-network simulations need to be able to address the influence of noise models on the effects of tensor-network truncation. This could provide a way to capture realistic features of near-term quantum circuits, such as limited entanglement and exploration of Hilbert space, in a way that benefits from the noise in these devices rather than being hindered by the typical expense of noise modeling.

Timeliness or maturity: With the rapid hardware development of GPU/TPU[8,9] based technologies due to AI and data science, classical simulation for quantum computing can be pushed to a different level after a decade of accelerated progress in computer technologies. With potential advances in adaptive basis truncation techniques and technology upgrades for the state-vector simulation, we believe the classical simulation can become a respectful partner for quantum computing at an intermediate scale to address quantum computing issues in the NISQ era and beyond, such as understanding performance and limitations of current and future quantum computers.

References

[1] J. Broz *et al.*, *DOE ASCR basic research needs in quantum computing and communications pre-workshop report*, (2023).

[2] Y. I. Manin, *A course in mathematical logic for mathematicians*, 2nd ed., Graduate Texts in Mathematics (Springer-Verlag, New York, NY, 2010).

[3] W. Kohn, *Nobel lecture: electronic structure of matter—wave functions and density functionals*, Rev. Mod. Phys. **71**, 1253 (1999).

[4] X. Xu *et al.*, *A herculean task: classical simulation of quantum computers*, (2023), arXiv:2302.08880 [quant-ph].

[5] J. Dalibard *et al.*, *Wave-function approach to dissipative processes in quantum optics*, Phys. Rev. Lett. **68**, 580 (1992).

[6] R. Islam *et al.*, *Onset of a quantum phase transition with a trapped ion quantum simulator*, Nat Commun **2**, 377 (2011).

[7] Y. Zhou *et al.*, *What limits the simulation of quantum computers?*, Phys. Rev. X **10**, 041038 (2020).

[8] A. G. M. Lewis *et al.*, *Large-scale distributed linear algebra with tensor processing units*, Proc. Natl. Acad. Sci. **119**, e2122762119 (2022).

[9] A. Morningstar *et al.*, *Simulation of quantum many-body dynamics with tensor processing units: Floquet prethermalization*, PRX Quantum **3**, 020331 (2022).

Title: Quantum algorithm development for simulating open quantum systems with energy applications

Authors: Fei Wang, fwang22@gmu.edu, Department of Chemistry & Biochemistry, George Mason University

Topic: applications, algorithms.
Charge transfer and exciton migration in electronics and photonics play a critical role in their performance regarding energy harvesting and conversion efficiency [1–3]. Computer simulations of these quantum dynamical processes can provide atomic level description and fine time resolution that are not always accessible through experiments [4,5]. As the quantum particles and excitations are embedded in a thermal bath, their accurate descriptions require the approaches of open quantum systems [6]. Although important progress has been made in developing classical algorithms for simulating such processes [7–9], the simulations on classical computers remain resource intensive due to the delocalized nature of quantum mechanics. Quantum computers, on the other hand, are naturally suited for quantum simulations. Therefore, it is beneficial to investigate the potential quantum speedup of using quantum computers to simulate open quantum system dynamics that underlies the fundamental process of many energy materials.

Challenge: Although there are approximate algorithms for the quantum computer simulation of open quantum system dynamics, such as using the Kraus operators for the Lindblad equation, numerically exact quantum algorithms for such simulations have not been developed. There are known classical algorithms that can produce numerically exact results, such as quasi-adiabatic path integral (QuAPI) [7], hierarchical equation of motion (HEOM) [8], and multi-configurational time-dependent Hartree (MCTDH) [9], but how much they can be adapted to implement on a quantum machine is unknown. In addition, the open quantum systems evolve under non-unitary operations. What are the feasible mathematical methods to convert non-unitary operators into unitary ones is an open question. Specially, the question to which scheme can lead to hardware efficient quantum circuits remains to be answered. A follow-up question would be how to use measurement to retrieve the relevant information related to the non-unitary operation after it is encoded in the unitary one.

Opportunity: As we are seeking practical applications of quantum computing, quantum computer simulation of open quantum systems remains an important area of investigation. Fundamentally, a quantum computer naturally eliminates the exponential scaling problem on a classical computer. Although there has not been a quantum counterpart of the successful classical algorithm outlined in the challenge, the classical algorithms can serve as starting points to build the quantum algorithm. In addition, the successful quantum algorithms applied to other areas such as variational quantum circuit [10], imaginary time evolution [11], and quantum signal processing [12] might be adapted and retooled for open quantum systems. There are great opportunities for interdisciplinary involvement in this endeavor of algorithm development, engaging mathematicians and computer scientists to prove the complexity of the problem and the degree of the quantum speedup. It also invites the codesign process between the hardware and software team.

Assessment: The success of the quantum algorithm will be capable of addressing important dynamical properties of open quantum systems such as the non-Markovian process and finite temperature effect. The successful algorithm will also take consideration of the symmetry in the system and produce hardware efficient quantum circuit that can be implemented on the NISQ devices. It will not have significant classical overhead. Gate counts and circuit depth are important metrics for a successful algorithm. The quantum speedup can be assessed based on the query model and the gate model.

Timeliness or maturity: With the number of qubits on quantum computers grow each year, it is the time to utilize the large dimension Hilbert space these qubits can provide to simulate open quantum systems. With inspirations from the algorithm development in neighboring fields, numerically exact quantum algorithm for simulating open quantum systems can be developed. The algorithm will have the capacity to treat multi-level and multi-site dynamics directly relevant to charge transfer processes in semiconducting devices and exciton migration in light harvesting complexes. Capable of conducting classically intractable simulations, quantum computers will offer critical insight into the functional mechanisms of the energy materials and inspire the rational design.

References

[1] A. E. Jailaubekov et al., *Hot Charge-Transfer Excitons Set the Time Limit for Charge Separation at Donor/Acceptor Interfaces in Organic Photovoltaics*, Nat. Mater. **12**, 66 (2013).
[2] T. Brixner, R. Hildner, J. Köhler, C. Lambert, and F. Würthner, *Exciton Transport in Molecular Aggregates - From Natural Antennas to Synthetic Chromophore Systems*, Adv. Energy Mater. **7**, 1700236 (2017).
[3] P. V. Kamat, *Boosting the Efficiency of Quantum Dot Sensitized Solar Cells through Modulation of Interfacial Charge Transfer*, Acc. Chem. Res. **45**, 1906 (2012).
[4] S. Kundu, R. Dani, and N. Makri, *Tight Inner Ring Architecture and Quantum Motion of Nuclei Enable Efficient Energy Transfer in Bacterial Light Harvesting*, Sci. Adv. **8**, eadd0023 (2022).
[5] R. Binder, D. Lauvergnat, and I. Burghardt, *Conformational Dynamics Guides Coherent Exciton Migration in Conjugated Polymer Materials: First-Principles Quantum Dynamical Study*, Phys. Rev. Lett. **120**, 227401 (2018).
[6] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems*, 1st ed. (Oxford University PressOxford, 2007).
[7] N. Makri, *Numerical Path Integral Techniques for Long Time Dynamics of Quantum Dissipative Systems*, J. Math. Phys. **36**, 2430 (1995).
[8] Y. Tanimura, *Numerically "Exact" Approach to Open Quantum Dynamics: The Hierarchical Equations of Motion (HEOM)*, J. Chem. Phys. **153**, 020901 (2020).
[9] M. Beck, *The Multiconfiguration Time-Dependent Hartree (MCTDH) Method: A Highly Efficient Algorithm for Propagating Wavepackets*, Phys. Rep. **324**, 1 (2000).
[10] M. Cerezo et al., *Variational Quantum Algorithms*, Nat. Rev. Phys. **3**, 625 (2021).
[11] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, *Variational Ansatz-Based Quantum Simulation of Imaginary Time Evolution*, Npj Quantum Inf. **5**, 75 (2019).
[12] G. H. Low and I. L. Chuang, *Optimal Hamiltonian Simulation by Quantum Signal Processing*, Phys. Rev. Lett. **118**, 010501 (2017).

Hamiltonian-Oriented Quantum Algorithm Design and Programming

Jiaqi Leng, Yuxiang Peng, and Xiaodi Wu[1]
University of Maryland, College Park

Topic: We propose a new paradigm of the design and programming of quantum algorithms based on the abstraction of Hamiltonian evolution, rather than the conventional quantum circuit model. Circuit-based digital abstraction has been successful in scaling up the design and implementation of modern classical computing chips, however, under the condition of the abundance of computing resources where correctness becomes a major issue for scalability. However, the limitation of computing resources is a major bottleneck for quantum computing in the near future. The current compilation pipeline from circuit-based quantum algorithms to pulse-level implementation on quantum hardware often incurs huge overheads. Moreover, the circuit abstraction also ignores the low-level but native programmability of underlying quantum physical devices and hence fails to fully leverage them. On the other side, quantum Hamiltonian evolution is a native abstraction for both low-level hardware control (programming and compilations) and many high-level applications (algorithms).

Thus, we propose to use quantum Hamiltonian evolution as the central object in end-to-end quantum application design: by designing Hamiltonian-based (i.e., to use Hamiltonian evolution as the first class object in programming) quantum algorithms and developing software stacks for programming and compiling quantum Hamiltonians natively to current or near-term quantum devices, we hope to significantly shorten the timeline towards demonstrating useful quantum applications with a full-stack toolchain.

Challenges: The conventional quantum circuit model, although with the aforementioned limitations for near-term applications, has long served as the well-recognized interface between theorists and experimentalists, which has greatly influenced the research agenda for each associated community to quantum computing. One major challenge in Hamiltonian-oriented algorithm design and programming is to identify such new Hamiltonian-based computational models, serving as an intermediate layer between applications and quantum hardware, which would be acceptable to both theorists and experimentalists. Ideally, we wish this abstract Hamiltonian layer allows efficient compilation to native operations on the quantum hardware. Meanwhile, this abstract Hamiltonian layer must be expressive enough to represent a rich class of existing and potentially new quantum algorithms. Achieving both requirements is challenging, which likely involves active collaborations between theorists and experimentalists, and will take iterations.

One foreseeable challenge from the application aspect is the design of quantum algorithms without quantum circuits. Although some quantum algorithms can be naturally formulated as Hamiltonian evolution directly, e.g., quantum adiabatic algorithm and continuous-time quantum walk, it remains an open question if we can convert other useful quantum routines (e.g., amplitude amplification, QFT, LCU, etc.) to Hamiltonian simulation in an intuitive and efficient manner. It is also equally challenging and exciting to discover new quantum algorithms directly based on Hamiltonian abstraction, although we might need to investigate unexplored application domains where computational tasks are more continuous in nature.

The challenge is also big in the development of software for controlling continuous-time quantum systems focusing on low-level programming of pulse shapes. Programming languages for intermediate-scale Hamiltonians are rarely developed, restricting the users' ability to design and implement Hamiltonian-oriented quantum algorithms for their computational tasks. For example, QuTiP, one of the state-of-the-art programming languages for continuous-time quantum systems, employs matrices to store data, which limits the system size that can be expressed. Compilation from Hamiltonian abstraction to pulse-level control could be much more sophisticated than the compilation for gate/circuit abstraction. Moreover, dealing with heterogeneous architectures of various quantum hardware due to their drastically different physics and pulse engineering techniques is also a big challenge.

Opportunities: Investigating the Hamiltonian abstraction is a concrete goal and a great opportunity for collaboration between theorists and experimentalists toward the purpose of hardware-efficient quantum application design. Some recent developments could serve as the starting point for this kind of research: e.g., Leng et al.[2]

[1]jiaqil@, ypeng15@, xiaodiwu@umd.edu.

[2]Quantum Hamiltonian Descent. `https://arxiv.org/abs/2303.01471`.

developed the so-called Quantum Hamiltonian descent (QHD), which is proposed as the genuine quantum counterpart of gradient descent, that can be implemented on the so-called Quantum Ising machine (QIM) through newly developed Hamiltonian embedding techniques. Here, QHD is an example of new quantum algorithms designed fully based on Hamiltonian abstraction and QIM is a candidate Hamiltonian modeling of quantum devices that allows minimal overhead in compilation to native machine instructions. Thanks to the QIM abstraction, QHD can be implemented nowadays on analog quantum simulators such as D-Wave and QuERA machines. Other Hamiltonian modelings of (e.g., bosonic or photonic) quantum devices are also worth investigating for hardware-efficient algorithm implementation.

In principle, any efficient quantum algorithm can be reduced to a (2-local) Hamiltonian simulation task because of the BQP-completeness of Hamiltonian simulation. However, this reduction comes with polynomial overhead (which can be huge in practice) and it usually results in extremely complicated quantum Hamiltonian as the final product, preventing efficient programming and implementation. We will need creativity in formulating existing quantum applications based on Hamiltonian abstraction (e.g., see an inspiring recent example [3]), and at the same time keep in mind that Hamiltonian abstraction might inspire novel quantum applications for domains whose nature is less discrete and combinatorial but more continuous and analytical.

A programming language and compiler design based on Hamiltonian abstraction will be critical to quickly deploy Hamiltonian-oriented quantum applications to various pulse-programming-enabled quantum platforms (e.g., IBM via OpenPulse, QuERA, and Rigetti via Braket pulse), which just become available very recently. An intuitive domain user experience, portability to various quantum platforms, efficiency, and overall productivity, will be important factors for the success of such programming infrastructure. A recently developed domain-specific programming language for quantum simulation called SIMUQ [4] could serve as a candidate framework for such infrastructure. In particular, the newly developed abstract analog instruction set (AAIS) abstraction provides a way to formally describe the native programmability of heterogeneous quantum devices while allowing efficient compilation based on Hamiltonian abstractions.

Many developments can be built on top of such an infrastructure: (1) benchmark sets for evaluating the performance of various applications and machines; or (2) toolchains that further facilitate domain applications from high-energy physics, quantum chemistry, or even the design of the quantum computer itself. Such an infrastructure will also significantly ease the collaborations between algorithm designers and hardware providers, and as a result, facilitate the codesign of both algorithms and hardware architectures.

Assessment: The realization of existing quantum algorithms based on Hamiltonian simulation using the aforementioned software tools will be the first step as proof of concept. Next, the emergence of Hamiltonian-based quantum algorithms with large-scale implementations for real-world applications will be a major milestone on our roadmap, marking the beginning of intense and regular applications of quantum computation in scientific and commercial domains.

To assess the software tools, we could evaluate the programming and compilation infrastructure on benchmarks based on metrics like run time, success rate, evolution fidelity, and energy consumption on real devices. User experience and overall productivity should be investigated through the evaluation of, e.g., #lines of code, coding times, or so with the language.

Timeliness/Maturity: Recent developments make pulse-level control of intermediate-scale quantum systems feasible over quite a few quantum platforms (e.g., superconducting, neutral atoms, and trapped ions), while scalable fault-tolerant quantum devices are still far from reach. We believe a promising pathway to deliver the first wave of useful quantum applications is through the direct programming of Hamiltonians.

Indeed, recent experiments have demonstrated the use of quantum simulators to exhibit quantum phenomena in an unprecedentedly feasible regime or to solve scaled-down versions of scientific problems from high-energy physics and quantum chemistry. This makes the Hamiltonian-oriented algorithm design both theoretically and practically feasible. Moreover, a lot of existing experimental demonstrations rely on the manual realization of simulation algorithms on the circuit level, which is an error-prone and tedious procedure, and incurs overheads in the final compiled pulse-level control instructions. This makes the development of Hamiltonian programming languages and compilations both necessary and timely.

[3] Quantum simulation of partial differential equations via Schrodingerisation. `https://arxiv.org/abs/2212.14703`.

[4] SIMUlation language for Quantum (SIMUQ). `https://pickspeng.github.io/SimuQ/`.

On the Road to Quantum Network: Vision, Challenges, Insights, and Opportunities

Wenji Wu*, Erhan Saglamyurek*, Ezra Kissel*, Hartmut Haffner+, Inder Monga*
wenji@lbl.gov, esaglamyurek@lbl.gov, kissel@es.net, hhaeffner@berkeley.edu, imonga@lbl.gov
* Lawrence Berkeley National Laboratory, + UC Berkeley

Topic: applications, algorithms, codesign and integration

Challenges: This position paper presents the QUANT-NET research team's vision on quantum network R&D from a network system perspective. A few challenges have been identified, which are grouped into two categories:

Toward practical deployment of quantum networks. Today, quantum networks are in their infancy. Like the Internet, quantum networks are expected to undergo different stages of research and development until they reach a level of practical functionality [1]. Existing quantum network testbeds and prototypes are typically implemented in laboratory experiments with limited functionality and in a tightly coupled manner. To move from laboratory experiments to practical deployment of quantum networks, a few challenges need to be addressed. *(1) Quantum network architectures and protocol stacks*. First, as the sizes and complexity of quantum networks increase, it is evident that quantum networks cannot be operated by manual control. The issues of quantum network abstraction, scaling of network architecture and protocols, and software automation are becoming increasingly important. In the OSI reference model, the communications between nodes are split into different abstraction layers. We believe that a similar layer network stack is likely constructed for quantum networks. The open issues are how to assign functions to each layer to realize quantum communication and networking. Second, entanglement is the fundamental building blocks of quantum networks [2, 3]. Key entanglement-related operations include entanglement generation & distribution, entanglement routing, entanglement swapping, and entanglement distillation. Therefore, how to build up a quantum network stack that can support various entanglement operations is critical. Open questions such as "*how to scale the protocol stack to handle thousands of, or even more, entanglement in parallel?"* need to be answered. Third, quantum networks aim for high-fidelity quantum state transport. We anticipate that incorporating quantum error correction into the quantum network protocol stack to achieve high-fidelity entangled links is critical [4]. *(2) Synchronization challenges.* Quantum resources in quantum networks must be fully coherent and synchronized [5, 6, 7]. Depending on quantum protocols to be implemented and underlying quantum technologies, synchronization in time, frequency, phase, or their combinations are required. High-fidelity Bell state measurements, an essential and important function in quantum networks, requires spectral, temporal, and polarization indistinguishability. As the sizes of quantum networks increase, quantum network control tends to be decentralized and distributed. Multiple synchronization references are likely to be deployed. Synchronization in such environments is challenging. While many synchronization technologies have already been developed and deployed in classical telecommunication networks, many of these still require significant improvements for quantum networks. For example, the frequency drifts in the QUANT-NET testbed must be controlled within a few kHz range. The temporal drifts in some quantum networks are required not to exceed a few picoseconds. These new synchronization requirements are orders of magnitude more demanding than before.

Quantum networking and computing co-design. History tells us that the Internet experienced significant growth after killer applications such as WWW and web browsers were developed. The significant growth of the Internet in turn revolutionized the computing industry. We envision that quantum networks will follow the same path as the Internet: quantum networking will be enabled by, and further stimulate, scalable quantum computing. Quantum networking and computing co-design will accelerate the development of both fields. The co-design considers the full quantum system stack from the top-level quantum applications to the bottom quantum networks. From the top, quantum applications provide insights into quantum computing paradigms and quantum state movement patterns to design and optimize the system. From the bottom, quantum networking issues and constraints need to be understood. In the middle, the system design effort needs to optimize and tradeoff to build the most optimal quantum computer considering both the top-down requirements and the bottom-up constraints. As a first step, two major challenges need to be addressed. *(1) Scaling quantum computers using modular quantum interconnects.* It is well believed that only when hundreds, even thousands, of logical qubits become available in a quantum computer can it do any useful computing work which is impossible by using classical means. Therefore, there is an urgent need to scale existing quantum computers to more powerful ones with many more physical qubits. To date, trapped-ion qubits and superconducting qubits are two leading quantum technologies for universal quantum computing, seeing considerable investments by industry. However, due to physical constraints (e.g., quantum noise) and control complexity, a latest single quantum processing unit (QPU) can only support a maximum of dozens of trapped ion qubits or a few hundred superconducting qubits although seminal work such as the multicore "QCCD" architecture has been proposed [8, 9]. To further scale quantum computers, the next generation quantum computers are likely to feature multiple QPUs, networked by modular quantum interconnects. Mechanisms and advanced technologies need to be developed. *(2) Quantum computing/networking programming interface.* Scalable quantum computing requires a quantum computing/networking programming interface in the software layer, allowing local applications to access remote quantum resources through quantum networks. Little work has been carried out in this field. In the 1980s, a group of researchers developed an application program interface for TCP/IP network communications called the socket interface [10]. The socket interface defines a variety of software functions or routines for the development of applications for TCP/IP networks, which significantly accelerate distributed computing and networking. We envision that an equivalent *quantum "socket" programming interface* should be developed for quantum computing and networking. Through this interface, quantum communication endpoints can be specified, and quantum QoS requirements such as fidelity and entanglement generation rate can be negotiated or specified.

Opportunities: (1) Building a few quantum network testbeds with a layered network architecture and functional protocol stack is key toward practical deployment of quantum networks. Because there are different types of quantum qubit technologies (e.g., trapped ion and superconducting), with each having its pros and cons, and the scales of quantum networks vary (e.g., LAN, WAN), different types of quantum network testbeds will likely be constructed to evaluate and compare concepts and technologies. Thus, different flavors of quantum network architectures and protocol stacks should be accordingly developed. To scale quantum networks to larger distances, advanced quantum repeater technologies are required. Quantum repeaters must be considered in the quantum network architecture and protocol design and development. In addition, quantum network modeling and simulation offers alternative approaches to study and research quantum network architectures and protocols without having physical networks. (2) Advanced technologies are required to ensure synchronization (time, frequency, phase, or their combinations) in quantum networks. New synchronization requirements and standards for quantum networks need to be established. (3) In terms of *scaling quantum computers using modular quantum interconnects*, advanced technologies such as efficient and high-fidelity (>99%) quantum light-matter interfaces, and superfast and low-loss switching fabric are required. (4) A new *quantum computing/network programming interface* needs to be developed.

Assessment: (1) Building a few quantum network testbeds with a layered network architecture and functional protocol stack is an initial success toward practical deployment of quantum networks. (2) The initial success for scaling quantum computers using modular quantum interconnects is to build either a trapped ion quantum computer with 100+ qubits, or a superconducting one with 1000+ qubits, featuring near full connectivity with >0.9999 quantum gates. Such a quantum computer can perform some practical quantum computing tasks. (3) A success criteria for the quantum computing/networking programming interface is that programmers can use this interface to develop applications with wide applicability.

Timeliness or maturity: (1) Building a few quantum network testbeds with a layered network architecture and functional protocol stack likely takes ~5 yrs. Practical deployment of quantum networks likely takes 10+ yrs. (2) Building a trapped ion quantum computer with 100+ qubits, or a superconducting computer with 1000+ qubits may take ~5 yrs. However, developing a full-fledged quantum computer that can solve real problems likely takes 15+ yrs. (3) While developing an initial version of a quantum computing/networking programming interface may take ~3-5 yrs, it is likely that this interface will evolve with advances in quantum network architectures and protocol stacks.

References

1. S. Wehner et al, *Science* **362**, 6412 (2018).
2. R. Horodecki et al, *Reviews of modern physics* **81**, 2 (2009).
3. W. Kozlowski et al, RFC **9340** (2023).
4. A. Ekert et al, *Phys. Rev. Lett*. **77**, 12 (1996)
5. A. Roulet et al, *Phys. Rev. Lett*. **121**, 6 (2018)
6. M. Lohe, *J. Phys. A: Math. Theor.* **43** 465301(2010).
7. R. Valivarthi, *J. Light. Technol. **40**, 23 (2022).*
8. D. Kielpinski et al, *Nature* **417**, 709 (2002).
9. J. Pino et al, arXiv: 2003.01293 (2020): 28-29.
10. W. Stevens et al, *ACM SIGCOMM. CCR* **20**, 2 (1990): 8-9.

Timing Synchronization and Control Electronics for Scalable Quantum Networks

Si Xie, Cristián Peña, Gustavo Cancelo, Leandro Stefanazzi, Panagiotis Spentzouris, Maria Spiropulu, Raju Valivarthi, Christina Wang, Neil Sinclair,

May 1, 2023

1 Introduction and Challenge

Scalable quantum networking is a critical scientific and technological milestone with a host of revolutionary and strategic applications spanning quantum computing, secure communications, and fundamental science. A key challenge that is required by all such applications is the precise synchronization of clocks across all nodes comprising the network [1, 2]. This is not only to ensure that remote nodes of the network operate in synchrony, but more crucially because quantum interference, which is required to produce entanglement, relies upon synchronization of photons. Clock synchronization to sub-picosecond level would enable many emerging technologies and improve protocols to be incorporated into quantum networks, such as those based on single photon detection [3], which offer orders of magnitude improvement in communication rates, and those utilizing spectrally distinct photons [4], which relaxes the requirements to achieve interference, as well as enable quantum-classical coexistence.

2 Opportunity

Recent rapid developments in high-speed radio frequency (RF) electronics spurred on by both industry needs as well as enhanced manufacturing processes have begun to standardize electronics in the tens of GHz approaching 100 GHz, thus enabling for the first time signal processing at the sub-picosecond level. In quantum network research and development, such emerging high-speed electronics have been recently employed to distribute and synchronize clocks to the few picosecond level across distances of several tens of miles, with clock signals coexisting in the same optical fiber as quantum signals. Concurrently, use of state of the art radio-frequency-system-on-a-chip (RFSoC) FPGAs have enabled sophisticated real-time control systems to be built and demonstrated for quantum network operations combining the fast and precise timing capabilities with flexible user-defined protocol controls [5]. The recent rapid development in these areas of electronics and control systems signals an opportunity to develop fully integrated control systems for quantum network operations, e.g. using emerging thin-film lithium niobate [6], protocol execution, error mitigation, and end-user quantum information delivery and handling.

3 Quantum Instrumentation Control Kit (QICK)

The QICK system was originally developed for Superconducting qubit experiments [7], but has since also been used to demonstrate entangled photon pair distribution and detection with time-bin photonic qubits [5]. QICK has been deployed on several different versions of the Xilinx RFSoC development boards and has the necessary versatility to adapt to many different types of quantum applications. Using the Xilinx RFSoC generation 3 UltraScale+ device, which can provide up to 16 output DACs at a speed of 10 Gsps and up to 16 input ADCs at 2.5 Gsps, the QICK system is an ideal platform for applications in quantum network controls. With a growing user base and rapidly developing firmware and software ecosystem, the QICK system is an ideal candidate to achieve the desired network-wide clock synchronization and to integrate the operation and controls of not only quantum networks, but also future hybrid architectures involving networked quantum computers and quantum sensors.

4 Assessment

Successful implementations of network-wide clock synchronization and integrated control systems include measurement and evaluation of metrics such as synchronization time resolution, synchronization and network operation stability, fidelity of key quantum network protocols, and ultimately the accuracy and success rate of quantum information delivery. In particular, the use of our control system is necessary to keep pace with the reduction of timing jitter of single photon detectors, which is reaching ps-levels for superconducting nanowires [8]. To further expand widespread quantum network usage towards the realization of the quantum internet, it is also prudent to minimize the cost of building such systems in order to facilitate the scalability of implementing quantum networks over a large geographical region and an increasing number of nodes.

5 Summary

A critical technological component of achieving state-of-the-art scalable quantum networks is clock synchronization and control systems. As high-speed electronics becomes further integrated into versatile control systems powered by increasingly powerful and customizable state-of-the-art FPGAs, it becomes possible and desirable to implement these key functions into the same device. The QICK system is one recent, and rapidly developing, highly successful example. Further progress and increasing success in this area of R&D will not only fulfill essential needs of operating quantum networks, but will also decrease the costs of deploying the necessary high-speed electronics by an order of magnitude compared to conventional solutions, thereby accelerating the realization of scalable quantum network systems. Finally, the QICK system will also find applications in other crucial networking technologies, ranging from quantum frequency conversion, classical-quantum coexistence, to protocol optimization as outlined in other position papers that have been submitted.

References

[1] R. Valivarthi et al., ”Picosecond synchronization system for quantum networks”, IEEE Journal of Lightwave Technology Vol 40 7668 (2022), doi: 10.1109/JLT.2022.3194860.

[2] K. Kapoor et al., ”Picosecond Synchronization System for the Distribution of Photon Pairs through a Fiber Link between Fermilab and Argonne National Laboratories”, IEEE Journal of Quantum Electronics, doi: 10.1109/JQE.2023.3240756.

[3] Lucamarini, Marco, et al. ”Overcoming the rate–distance limit of quantum key distribution without quantum repeaters.” Nature 557.7705 (2018): 400-403.

[4] Vittorini, G., et al. ”Entanglement of distinguishable quantum memories.” Physical Review A 90.4 (2014): 040302.

[5] S. Xie et al., ”Entangled Photon Pair Source Demonstrator using the Quantum Instrumentation Control Kit System”, arxiv:2304.01190, submitted to IEEE Journal of Quantum Electronics (2023).

[6] Zhu, Di, et al. ”Integrated photonics on thin-film lithium niobate.” Advances in Optics and Photonics 13.2 (2021): 242-352.

[7] L. Stefanazzi et al., “The QICK (Quantum Instrumentation Control Kit): Readout and control for qubits and detectors,” Review of Scientific Instruments Vol 93, 04470 (2022).

[8] Korzh, Boris, et al. ”Demonstration of sub-3 ps temporal resolution with a superconducting nanowire single-photon detector.” Nature Photonics 14.4 (2020): 250-255.

Hybrid Quantum Networks: Modeling and Optimization

Yufeng Xin
RENCI, University of North Carolina
Chapel Hill, NC, USA
yxin@renci.org

Inder Monga, Liang Zhang
ESnet, LBNL
Berkley, CA, USA
imonga@es.net, lzhang9@es.net

1 Topic: Quantum networking: architecture, models, and algorithms

2 Challenges

The well-received Quantum Internet evolution roadmap indicates that quantum network will be of a hybrid satellite-fiber quantum network system in the foreseen future [1]. With the ever expanding capabilities of quantum computing, the roles of classical and quantum communication channels (in-band or out-band) in the hybrid system will need to evolve accordingly in order to maximize the benefits. However, current research on quantum networks treats classical and quantum networks as completely separate networks, with the classical channel serving as the control plane for the quantum network, as proposed by several SDN approaches [2].

The different architectural possibilities of the hybrid satellite-fiber quantum networks pose unprecedented challenges in three areas: the control of the hybrid satellite and fiber quantum links, the coexistence of traditional and quantum networks, and the fusion of classical network models and quantum characteristics.

A significant challenge for this hybrid system is that the limited quantum channel capacity and stability continue to be major obstacles, despite recent progress in QKD (Quantum Key Distribution) and entanglement hardware that has demonstrated promising results for transporting critical information over distances [3, 4]. This calls for innovations in network modeling and optimal control in order to accelerate the progress towards the grand quantum Internet vision.

3 Opportunity

3.1 Hybrid Network Architecture and Use Cases

In such a system, quantum nodes are strategically deployed in chosen ground and satellite locations. There could be multiple variants of the hybrid network architecture.

- A classical control plane for the quantum network.
- A quantum control plane for the classical Internet. A QKD network can be extended to assume the role of a secure control plane for the classical network at scale.
- A hybrid transport network with quantum segments and traditional Internet segments.
- Quantum Network. Due to the wide coverage, satellite quantum nodes can play multiple roles, for example, Trusted relay, Quantum repeater, and Quantum memory. When the network size scales up, the satellite quantum nodes can also play the role of cluster head for partitioned control of the fiber quantum network on the ground.

3.2 Network Modeling, Optimization, and Topology Control

The burgeoning number of papers in quantum network in most recent years focused on the routing problem, most of which are customized or heuristic extensions of existing path and network flow algorithms. With the limited number of qubits the quantum device can generate, a quantum entanglement network is a capacitated network consisting of unreliable links due to the fast decoherence phenomenon. Repeaters are imperative to extend the range of the network and multipartite entanglement may become the new normal rather than the traditional pairwise communication.

- Model. Adequate graph theoretical models, especially hypergraph models and probabilistic graph models, are needed to abstract out the quantum physical layer.
- Optimization. The network optimization problem space in quantum network is much bigger than that in the traditional network. In addition to the fundamental network design and traffic engineering (TE) problems, quantum node and repeater location, network partition, cross-layer, and fault tolerance are all prominent problems in various pure or hybrid network architectures. These are inherently NP hard. It would be intriguing if the quantum computing advancement in solving intractable optimization problems could be readily leveraged.
- Control. Difficulty in entanglement generation and low fidelity quantum links implies efficient stochastic decision models and optimal control algorithms are needed for quantum link formation and topology control.

4 Assessment

Simulation and emulation informed by the real testbed implementations will remain the main pipeline to advance and measure the quantum network solutions. Network throughput, reliability, and latency will be the key performance metrics. Network security concern may be completely eliminated.

5 Timeliness or maturity

The substantial qubit capacity increase by the quantum computer and breakthroughs in photonic, quantum memory and repeater devices foreshadow a percolation point of at-scale quantum network that may be reached sooner than previously expected.

References

[1] S. Wehner, D. Elkouss, and R. Hanson, "Quantum internet: A vision for the road ahead," *Science*, vol. 362, no. 6412, p. eaam9288, 2018. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.aam9288

[2] D. Awschalom, "From long-distance entanglement to building a nationwide quantum internet: Report of the doe quantum internet blueprint workshop," Brookhaven National Lab.(BNL), Upton, NY (United States), Tech. Rep., 2020.

[3] Z. Yang, M. Zolanvari, and R. Jain, "A survey of important issues in quantum computing and communications," *IEEE Communications Surveys & Tutorials*, 2023.

[4] C. Qiao, Y. Zhao, G. Zhao, and H. Xu, "Quantum data networking for distributed quantum computing: Opportunities and challenges," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2022, pp. 1–6.

Algorithm and hardware designs with native global quantum gates

Bin Yan[1,*], and Nikolai Sinitsyn[1]
[1] Los Alamos National Laboratory, Los Alamos, New Mexico 87545
[*] Representing author: byan@lanl.gov, (505) 665-5012

Topic: Quantum computing models, algorithms and codesign

Challenge: In recent years, the field of quantum computing has undergone exponential progress. Notably, the current generation of Noisy Intermediate-Scale Quantum (NISQ) processors have surpassed classical counterparts and exhibited quantum supremacy on certain computational tasks. Nevertheless, the ultimate goal of constructing fault-tolerant quantum computers still remains a formidable challenge that has yet to be overcome. The quantum volumes of current NISQ processors remain significantly constrained by their noise levels, which hinders their potential to solve complex real-world scientific problems.

Quantum circuit designs commonly rely on universal sets of quantum gates, which typically consist of single and two-qubit quantum gates, such as the CNOT gate. However, certain algorithmic components, when transcribed onto a quantum circuit, may demand a vast number of few-qubit gates from the universal set. This presents a challenge, as these few-qubit gates may not always be the most natural operation that can be directly executed on the physical hardware, necessitating an additional compilation step.

The two aforementioned layers of compilation inevitably impose an unnecessary burden on the overall algorithm execution, often surpassing the limited quantum volume of NISQ processors. Consequently, a pressing challenge is to explore viable solutions to circumvent this limitation and elevate the functional capabilities of existing NISQ devices for practical applications.

Opportunity: Many quantum hardware devices have the ability to execute global quantum gates, which entail gate operations that involve multiple qubits simultaneously. For instance, ion trap qubits interact via the Mølmer-Sørensen interaction (Ising-type interactions), allowing for the creation of an untargeted unitary across the entire array of qubits through a global pulse control [1, 2]. In contrast, implementing targeted few-qubit operations is generally more challenging. Integrating these global operations as a native component into the algorithm design can significantly alleviate the quantum volume required for the computation.

Let us take an example to see what global quantum gates can offer: Grover's algorithm provides a quadratic speedup for the unsorted database search problem. but it assumes a global, problem-specific, oracle operation among all qubits. The implementation of this operation using few-qubit gates may be so complex that it ultimately erodes the quantum speedup. In a recent work [3], a novel solution to the number partition problem, which is an NP-hard problem with substantial practical relevance, was proposed by employing the Grover algorithm as a critical building block. Notably, the requisite oracle operation was implemented by applying a single global control to all the qubits interacting via the Hamiltonian

$$H = \sum_{i}^{n} s_i \sigma_z^i \sigma_z,$$

where s_i are the integer numbers in the number partition problem, and σ_z is the spin-$1/2$ Pauli matrix. The Hamiltonian is comprised of solely two-body star-like interactions, which can be conveniently realized using platforms such as Rydberg-atom or cavity-QED.

What is the complexity, or the scaling of the necessary resources to implement the aforementioned oracle? In [4], it was demonstrated that the oracle can be alternatively achieved by driving the same Hamiltonian through a (quasi-)adiabatic evolution. Notably, the runtime scales logarithmically with the problem size, preserving the quadratic speedup of the Grover algorithm. Furthermore, this adiabatic oracle is topologically protected, rendering it robust against small imperfections in the control pulses and interaction parameters. This example illustrates how an algorithm and hardware co-design that employs native global quantum gates can amplify the capabilities of current NISQ platforms and advance practical scientific applications of quantum computation.

Recent developments in both theory and experimental techniques have presented a significant opportunity to delve into this direction. Some of the challenges that need to be tackled include characterizing the native set of global quantum gates that a given physical processor can execute, conversion between standard two-qubit gates and global quantum gates, algorithm design tailored to specific physical architectures, and more.

Assessment of success is multi-faceted. On the theoretical front, success would entail deepening our understanding of the structure of global quantum gates. Additionally, the complexity of implementing global gates, such as scaling of the operation time with respect to system size, needs to be assessed. Moreover, it is crucial to evaluate the performance of algorithms using global quantum gates compared to conventional approach using few-qubit gates. As for experimental evaluations, the implementations of the global quantum gates must be assessed using various metrics and benchmarking techniques, including fidelity, operation time, etc.

Maturity: The concept of utilizing global quantum gates is not a new one; few-qubit gates have never been the only option for the universal set of quantum gates. While they are convenient for expressing abstract quantum circuits, they may not always be optimal for physical implementations. However, it is only recently, with advancement in experimental techniques leading to higher fidelity of global unitary operations, that research on algorithms with global quantum gates has become practically relevant. In the industry, corporations are investing to leverage the opportunity of global gates as well, such as QuEra and Infleqtion (ColdQuanta). This has led to the emergence of a new field of study on global quantum gates. To name a few recent references in addition to the work cited above, global ease gates on ion trap qubits was considered in [5], constant cost realization of global entangling gates was studied in [6], compilation of circuits with 2-qubit gates into global gates was investigated in [7]. In our view, investing in this area is a timely effort. We therefore propose it as a topic for discussion in the ASCR workshop on quantum computing and networking.

References:

[1] S. Debnath et al. “Demonstration of a small programmable quantum computer with atomic qubits”. In: *Nature* 536.7614 (Aug. 2016), pp. 63–66.

[2] C. Figgatt et al. “Parallel entangling operations on a universal ion-trap quantum computer”. In: *Nature* 572.7769 (Aug. 2019), pp. 368–372.

[3] G. Anikeeva et al. “Number Partitioning With Grover’s Algorithm in Central Spin Systems”. In: *PRX Quantum* 2.2 (May 2021), p. 020319.

[4] B. Yan and N. A. Sinitsyn. “Adiabatic oracle for Grover’s algorithm”. In: *arXiv: 2207.05665 (quant-ph)* (July 2022).

[5] N. Grzesiak et al. “Efficient quantum programming using EASE gates on a trapped-ion quantum computer”. In: *Quantum* 6.634 (Jan. 2022).

[6] S. Bravyi et al. “Constant-Cost Implementations of Clifford Operations and Multiply-Controlled Gates Using Global Interactions”. In: *Phys. Rev. Lett.* 129.23 (Nov. 2022), p. 230501.

[7] J. van de Wetering. “Constructing quantum circuits with global gates”. In: *New J. Phys.* 23.4 (Apr. 2021), p. 043015.

Title: Measurement-based Quantum Algorithms for Near-term Quantum Hardware

Authors: Kübra Yeter-Aydeniz (MITRE Corporation, kyeteraydeniz@mitre.org);
George Siopsis (University of Tennessee, Knoxville, siopsis@tennessee.edu)

Topic: Applications, Computing and Programming Models, Algorithms, Error Correction and Mitigation

Challenge

Over the last three decades, quantum algorithms have been developed to solve various challenging problems. The majority of these algorithms are based on the circuit-based quantum computation (CBQC) model in which the quantum algorithm is implemented by applying a series of single- and two-qubit unitary operators followed by readout measurements at the end of the computation. Most of these algorithms have been tested on current quantum hardware. However, there is a plethora of quantum computing models. Measurement-based quantum computing (MBQC) (also known as one-way quantum computing) [1] is a model in which a large, entangled state is prepared initially, and computation is realized through single-qubit measurements on a subset of the qubits comprising the entangled state. The unmeasured qubits undergo the desired evolution controlled by the measurements. MBQC is *advantageous* compared to CBQC since generating a large, entangled state is easier than implementing the large number of unitary gates needed for CBQC to be performed, and this process does not change depending on the computation. In addition to this, fault-tolerant quantum computing requires quantum error correction (QEC) and most of the QEC schemes are based on consecutive measurements, thus relying on the same principles as MBQC. Nevertheless, the research area of MBQC is *underexplored*.

Challenge #1: It was not until recently that MB versions of some fundamental hybrid quantum-classical algorithms such as MB-VQE (Variational Quantum Eigensolver) [2], and MB-QAOA (Quantum Approximate Optimization Algorithm) [3] have been proposed. The research area of MB quantum algorithms and possible advantages they can bring is *underexplored*.

Challenge #2: There are plenty of metrics developed to quantify the characteristics and development of the CBQC model, such as randomized benchmarking, quantum volume, and cycle benchmarking. However, the MB counterpart of these metrics was not proposed until very recently [4]. The research area of MBQC metrics and benchmarks is *underdeveloped*.

Challenge #3: Similar to algorithms and metrics development, software that enables MBQC is scarce compared to CBQC software availability. Two example MB software are MCBeth [5] and Paddle Quantum. MB software development is in its *infancy*.

Opportunity

Challenge #1: Development of MB correspondent of hybrid quantum-classical CBQC algorithms, studying the resource overhead of these algorithms on near-term quantum hardware, and the effect of quantum hardware noise during implementation of these algorithms to solve real-world application problems.

Challenge #2: Development of metrics and benchmarks for the MBQC model and testing these metrics and benchmarks on existing quantum hardware.

Challenge #3: Development of open-source MB software with simulator and hardware implementation capability. Existence and availability of these software tools is also critical for addressing challenges #1 and #2.

Assessment

Challenge #1: We define success as performing useful computation with less error and less resource overhead compared to quantum algorithms that are based on the CBQC model or other computational models, such as quantum annealing. To this end, developed algorithms need to be tested in many application areas, such as solving challenging chemistry and physics simulation problems, optimization in real-world application areas. These applications should be followed by resource estimation studied for larger scale problems that are beyond today's quantum hardware capabilities.

Challenge #2: Initial step towards success in this area would be development of metrics that specifically address characteristics of MBQC. These metrics need to be tested on current quantum hardware in various quantum hardware architectures, such as superconducting, and photonics quantum hardware. Development of some of the quantum hardware features, such as mid-circuit measurement, is critical for realization of the MBQC model. Similar to CBQC, development of application specific benchmarks in MBQC is also critical to track progress in the field.

Challenge #3: Success in software development would be the ability of seamless transition from other computing models to MBQC, providing implementation in various quantum hardware, availability of optimization tools for resource reduction, as widely adapted by the scientific user community.

Timeliness and maturity

One of the most promising large-scale, fault-tolerant quantum hardware architecture is photonics which is the natural platform for MBQC. Recent developments in photonic quantum hardware are the biggest motivation for studying MBQC. For example, engineering of large cluster states with continuous variables, and development of programmable photonics quantum hardware. Success in MBQC can pave the way to fault-tolerant, large scale quantum computing.

References

[1] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest. *Measurement-based quantum computation*. Nature Physics 5, no. 1 (2009): 19-26.

[2] R. R. Ferguson, L. Dellantonio, A. Al Balushi, K. Jansen, W. Dür, and C. A. Muschik. "Measurement-based variational quantum eigensolver." Physical Review Letters 126, no. 22 (2021): 220501.

[3] M. Proietti, F. Cerocchi, and M. Dispenza. "Native measurement-based quantum approximate optimization algorithm applied to the Max K-Cut problem." Physical Review A 106, no. 2 (2022): 022437.

[4] Y. Zhang, D. Niu, A. Shabani, and H. Shapourian. "Quantum Volume for Photonic Quantum Processors." Physical Review Letters 130, no. 11 (2023): 110602.

[5] A. Evans, S. Omonije, R. Soulé, and R. Rand. "MCBeth: A Measurement Based Quantum Programming Language." arXiv preprint arXiv:2204.10784 (2022).

Challenges and Opportunities in Quantum Networking for Future Distributed Quantum Processing: New Perspective by Quantum Wrapper Networking

Authors: *S. J. Ben Yoo[1], Prem Kumar[2], Gregory S. Kanter[2], Sandeep K. Singh[1]; email: sbyoo@ucdavis.edu*
Organizations: *[1]University of California, Davis; [2]Northwestern University*
Position Paper in Response to: *ASCR Basic Research Needs in Quantum Computing and Networking*
Topics: *Computing and programming models; Codesign and integration across the quantum computing and networking stacks*

1. Challenges

The prospect of a quantum internet [1], [2] with quantum entanglement between any points on the Earth's surface offers truly fascinating possibilities. When quantum computers become available and commercially viable, then they will be able to solve problems that are exponentially difficult for classical computers. The imminent question is how to interconnect distributed quantum nodes and network devices having quantum state processing capability and manage a quantum network infrastructure for secure communication and distributed computing. A quantum network could be built on existing fiber-optic infrastructure [3], [4], however, despite some progress in developing quantum computing, networking, and related device technologies, the quantum internet is far from reality, as the fundamental properties of quantum mechanics forbid amplification, measurements or monitoring required for transportation, network control and management. Here, the fragile qubits need to be routed and switched while guaranteeing their integrity with some level of assurance in ***Quality of Transmission (QoT)*** and ***Quality of Entanglement (QoE).*** Furthermore, when such a quantum networking technology becomes available, our past experience tells us that **successful transitions** in networking technologies would require a strategy for **seamless upgrades** from today's classical networks to the new quantum networks with **interoperability** and possible **co-existence** of the two networks.

Many **new challenging questions** arise in considering the development and deployment of quantum networks: *(1) How do we place a **control plane** on quantum networks? (2) How do we **manage** quantum networks? (3) How do we **monitor the performance** of quantum networks? (3) How do we **switch, route** and achieve **end-to-end transportation** in quantum networks? (4) How do we **stabilize polarization, and synchronize payload and header** during their transport? (5) How do we **codesign and integrate quantum protocols and algorithms in software stack**? (6) Do we have **simulation tools and experimental testbeds** to design, simulate, and operate quantum networks with quantum devices before actual deployment? (7) Once we somehow manage to find solutions to all of the above, how do **we interoperate and seamlessly upgrade from today's networks to future quantum networks**?*

2. Opportunities

The authors of this Position Paper recently invented **Quantum Wrapper (QW) Networking**[5] [6]technology inspired by the architecture and the protocol of the Optical-Label Switching networks [7] In QW networks, the Quantum Wrapper is composed of the **QW Header** and the **QW Tail** in the form of classical bits to contain information pertaining to routing, multiplexing, timing, format, priority, etc. These classical bits will 'wrap' (lead and follow) the **quantum payload (qubits)** to facilitate end-to-end transport and switching of the quantum payload without reading or altering the quantum data payload until it reaches the qubit receiver. This QW networking method potentially offers the following opportunities:

1. QW networking [5] [6] can be deployed independently or in **co-existence with classical networks** while **supporting full interoperability and backward compatibility** with classical networks.
2. QW networking is a **transparent** optical networking technology utilizing classical quantum wrappers without reading the quantum data payload (qubits). Because of this transparency, (a) QW networking offers **multi-user and multi-quantum-entanglement-state distribution functionalities**, and (b) QW payload (qubit) can be of **any protocol and format** [5] [6].
3. QW networking **exploits much of the existing control plane protocols** to allow backward compatibility and seamless upgrades from today's networks to the future quantum internet. As already demonstrated in optical label switching, QW network enables a **software-defined-network (SDN),** where an out-of-band data communication channel (DCC) can be used to communicate control information between the quantum networking nodes and a network control and management (NC&M) system over a Data Communication Network (DCN) in a centralized NC&M system[8], while QWs themselves can offer distributed control of QW network nodes, similarly to IP networks. The QW

network control and management achieves interoperability and **backward-compatibility with existing or legacy telecom protocols** including Ethernet, OTN, MPLS, etc.

4. QW networking utilizes classical bits in the QWs to conduct **optical performance monitoring** to infer Signal-to-Noise-Ratio (SNR), QoT, Dispersion, Polarization Mode Dispersion (PMD), etc., without touching the qubits. Here the QWs at relatively low speed can be used as a supervisory classical channel to monitor the inferred quality of the copropagating quantum channels [9] used for, for example, distributing quantum information, including entanglement.
5. QW networking **does not require strict synchronization between the datagrams**, and thus facilitates development and deployment of the QW node systems (i.e. QW Switching Routers). On the other hand, QWs themselves help synchronization and polarization stabilization for QW payloads (qubits) for the qubit receivers due to the intimate integration of the QW header/tail with the QW payload.
6. QW networking can **help develop quantum network TCP/IP** due to the QW mechanisms that can be correlated with the qubit receivers.
7. QW networking mechanism **facilitates codesign and integration of compute and network stack**, which can be fully automated and software controlled.
8. QW network experimental testbeds can **facilitate the development of realistic transmission and quantum impairment models**, as well as simulation tools under different noise settings. Creating tools and mechanisms for communicating between different layer interfaces is necessary. The testbeds can address understanding of key impairments as well as the utility of indirect monitoring techniques for the required quantum bit error rates (qBER), and eventually the seamless upgrade scenario studies.

3. Assessment

Assessment by benchmarking various quantum networking techniques including QW would be important, by both simulation and experimental methods. In addition to classical networking systems' metrics of evaluation (such as bit error rate, signal to noise ratio, throughput, goodput, latency, etc.), the state of quantum bits requires assessment of coincidence count, fidelity, qSNR, qBER, Bell state measurement, coherence, etc. It is essential that systems and processes external to the quantum data plane have visibility into the state of the quantum network resources (e.g., wavelengths, quantum repeaters, etc.), and be able to assess the quality (via performance monitoring) of quantum data flowing within that data plane.

4. Timeliness

Today, with recent advances in quantum computing technologies, quantum networking has emerged as important area of research. There has been significant progress made in quantum network architectural design [3]–[6], [10]–[14], entanglement distribution [4], [13], [14], repeaters [12], and many other technologies, leading to development of experimental testbeds [4], [12]–[14], simulation tools, and detailed theories pertaining to quantum networks. Time is ripe for us to pursue quantum networking at full speed.

5. References

[1] H. J. Kimble, "The quantum internet," *Nature*. 2008.

[2] S. Wehner, et al, "Quantum internet: A vision for the road ahead," *Science*, vol. 362, no. 6412 (2018).

[3] M. Pant, et al, "Routing entanglement in the quantum internet," *npj Quantum Inf*, vol. 5, no. 1, 2019.

[4] M. Pompili, et al, "Experimental demonstration of entanglement delivery using a quantum network stack," *npj Quantum Inf*, vol. 8, no. 1, 2022.

[5] S. J. Ben Yoo and P. Kumar, "Quantum Wrapper Networking," in *IEEE Photonics Conference (IPC)*. Paper ME 1.3, 2021.

[6] S. J. Ben Yoo, et al, "Quantum Wrapper Networking." : https://arxiv.org/submit/4870190.

[7] S. J. B. Yoo, "Optical-label switching, MPLS, MPLambdaS and GMPLS," SPIE APOC 2001.

[8] S. J. B. Yoo, "Optical packet and burst switching technologies for the future photonic Internet," *Journal of Lightwave Technology*, vol. 24, no. 12, 2006.

[9] D. J. Geisler, et al, "Experimental demonstration of flexible bandwidth networking with real-time impairment awareness," *Opt Express*, vol. 19, no. 26, 2011.

[10] S. L. N. Hermans, et al, "Qubit teleportation between non-neighbouring nodes in a quantum network," *Nature*, vol. 605, no. 7911, 2022.

[11] S. Muralidharan, et al, "Optimal architectures for long distance quantum communication," *Sci Rep*, vol. 6, no. 1, 2016.

[12] M. K. Bhaskar, et al "Experimental demonstration of memory-enhanced quantum communication," *Nature*, vol. 580, no. 7801, 2020.

[13] J. Chung, et al, "Illinois Express Quantum Network (IEQNET): metropolitan-scale experimental quantum networking over deployed optical fiber," in *SPIE*. Quantum Information Science, Sensing, and Computation XIII. 2021.

[14] M. B. On, et al, "Demonstration of Quantum Channel Monitoring via Quantum Wrappers," in *Optical Fiber Communication Conference (OFC) 2023*, Paper Tu3H.4.

Enabling Multi-level Parallelism in Heterogeneous Quantum-Classical Programming Models

Akihiro Hayashi* Austin Adams* Jeffrey Young* Eugene Dumitrescu† Vivek Sarkar* Thomas M. Conte*
*Georgia Institute of Technology †Oak Ridge National Laboratory
Email: {ahayashi,aja,jyoung9,vsarkar,conte}@gatech.edu, dumitrescuef@ornl.gov

Topic: **Quantum applications and compilation**; We discuss our current and planned enhancements to the QCOR [1] programming system, which is based on the C++ programming language and is retargetable to different types of quantum computers, including 1) physical quantum hardware [2] and 2) quantum simulators [3] on conventional systems. Our enhancements are focused on increasing parallelism in QCOR programs and future quantum-classical backends.

Challenge: There are multiple levels of parallelism in quantum-classical hybrid algorithms, but existing quantum-classical programming models and these backends typically do not effectively support them. As a motivating example, we use Shor's algorithm. Algorithm 1 is a pseudo-parallel version of Shor's algorihm, where SHOR is a quantum-classical task that invokes the period-finding quantum kernel (SHORKERNEL) to estimate exponent r and SHOR can be called multiple times until one or more (non-)trivial divisors are found or the entire search space is explored. Also, **async** represents parallel task creation and execution and **foreach** represents parallel loop creation and execution. Assuming these constructs are specified users or compilers, one possibility of accelerating this algorithm is to run multiple instances of SHOR in parallel. Furthermore, since it can require multiple shots to find r, it would be also possible to further parallelize the shot loop in SHOR (Line 11). Finally, if the SHORKERNEL is executed on a simulator, there is a massive amount of parallelism as in [3]–[6].

In general, we identify the following multiple levels of parallelism in quantum-classical programs:

- Task level parallelism: multiple independent classical tasks that can include quantum kernels are executed in parallel.
- Shot level parallelism: multiple independent shots are executed in parallel.
- Inner simulator level parallelism: quantum simulators, including state vector and tensor network simulators such as [3]–[6], are typically parallelized using OpenMP, CUDA, and the Eigen library to utilize a massive amount of parallelism on CPUs and/or GPUs.

Therefore, we believe that exploiting different levels of parallelism in quantum-classical programming models on conventional systems will 1) accelerate the development of a quantum-classical algorithm, and 2) facilitate porting an existing heterogeneous algorithm to a quantum-classical one. However, there is no such quantum-classical programming model that aims to exploit the full capability of conventional systems.

Opportunity: We focus on enhancing the QCOR programming system as it is one of the state-of-the-art quantum-classical programming models developed at ORNL. Since QCOR is primarily written in C++, we look to enable user-level multi-threading in QCOR in a way that is acceptable to both QCOR and C++ programmers. For QCOR programmers, our goal is to minimize modifications to the code required for enabling multi-threading. For C++ programmers, our goal is to provide a threading interface that is natural to use. To that end, we leverage C++'s standard threading constructs (`std::thread` and `std::async`). However, in terms of general applicability, our discussions should apply to other parallel programming systems for C++, such as OpenMP [7], Kokkos [8], and RAJA [9]. Also, to implement parallel-aware backends, we identify portions in QCOR that can possibly inhibit user-level multi-threading. Essentially, these cases are focused on identifying potential sources of data races when multi-threading is added.

We have made preliminary modifications to QCOR and our preliminary paper [10] shows that enabling user-level multi-threading gives us performance improvements over the conventional baseline version in which each kernel is still executed by multiple threads, but is executed one-by-one. We believe that there are further group collaboration opportunities with teams working on quantum-classical algorithms to extend and advance this work for larger multi-threaded applications and test cases.

Assessment: We have recently measured the effectiveness of our enhanced programming model and its runtime system by mainly measuring end-to-end performance, and we provide results in Figure 1 from [10] showing strong scalability of two Shor's kernels with the one-by-one and the parallel approaches on a 12-core, 24-thread AMD Ryzen9 3900X CPU platform. The numbers are relative performance improvements over the single-threaded one-by-one execution. While both approaches show good scalability, the parallel version always outperforms the baseline, which indicates that two-level parallelization (concurrently running two quantum kernels with $N/2$-threads) is better than one-level parallelization (running kernels with N-threads one after the other) in a certain case.

In future work, we look to use different quantum-classical algorithms such as VQE to demonstrate that specifying multi-level parallelism accelerates prototyping and developing quantum-classical programs on conventional systems.

Timelines and maturity: We believe this multi-threading design for heterogeneous quantum-classical programming models will open up an opportunity for rapidly prototyping and developing quantum-classical programs on conventional systems with

Algorithm 1 Parallel Shor's Algorithm (Pseudocode)

Input: N: A natural number to be factorized.
Output: A non-trivial divisor(s) of N.

```
procedure MAIN(N)
    repeat
        a ← random(1, N);                                   ▷ 1 < a < N
        K ← gcd(a, N);
        if K == 1 then
            async SHOR(N, a);
        else
            return K
    until a divisor(s) is found or explored all
procedure SHOR(N, a)
    foreach s = 1, ..., nShots do
        r_s ← SHORKERNEL(N, a)
    r ← r_1, ..., r_s                    ▷ Estimate r from the measurements
    if r mod 2 ≡ 1 or a^r mod N ≡ −1 then
        return ϕ;
    else
        return gcd(a^r/2 ± 1, N);
```

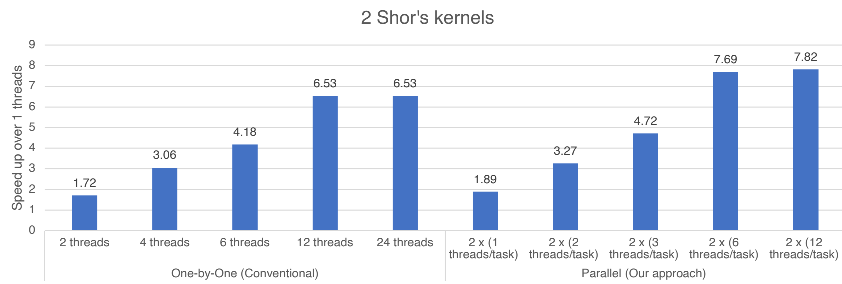

Fig. 1: Scalability of the one-by-one and the parallel approaches: two SHOR(N=7, a=2) from Algorithm 1

CPUs and GPUs. Since QCOR has emerged as one of the state-of-the-art generic and heterogeneous programming models, we anticipate that extensions to this initial design would be well utilized by researchers across the DoE, industry, and academia for longer-term explorations of heterogeneous programming systems for future quantum-classical systems.

REFERENCES

[1] T. M. Mintz, A. J. McCaskey, E. F. Dumitrescu, S. V. Moore, S. Powers, and P. Lougovski, "Qcor: A language extension specification for the heterogeneous quantum-classical model of computation," *J. Emerg. Technol. Comput. Syst.*, vol. 16, no. 2, mar 2020. [Online]. Available: https://doi.org/10.1145/3380964

[2] A. Adams, E. Pinto, J. Young, C. Herold, A. McCaskey, E. Dumitrescu, and T. M. Conte, "Enabling a programming environment for an experimental ion trap quantum testbed," in *2021 International Conference on Rebooting Computing (ICRC)*. Los Alamitos, CA, USA: IEEE Computer Society, dec 2021, pp. 14–23. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/ICRC53822.2021.00014

[3] V. Gheorghiu, "Quantum++: A modern c++ quantum computing library," *PLOS ONE*, vol. 13, no. 12, pp. 1–27, 12 2018. [Online]. Available: https://doi.org/10.1371/journal.pone.0208073

[4] Y. Suzuki, Y. Kawase, Y. Masumura, Y. Hiraga, M. Nakadai, J. Chen, K. M. Nakanishi, K. Mitarai, R. Imai, S. Tamiya, T. Yamamoto, T. Yan, T. Kawakubo, Y. O. Nakagawa, Y. Ibe, Y. Zhang, H. Yamashita, H. Yoshimura, A. Hayashi, and K. Fujii, "Qulacs: a fast and versatile quantum circuit simulator for research purpose," *Quantum*, vol. 5, p. 559, Oct. 2021. [Online]. Available: https://doi.org/10.22331/q-2021-10-06-559

[5] T. Vincent, L. J. O'Riordan, M. Andrenkov, J. Brown, N. Killoran, H. Qi, and I. Dhand, "Jet: Fast quantum circuit simulations with parallel task-based tensor-network contraction," *Quantum*, vol. 6, p. 709, May 2022. [Online]. Available: https://doi.org/10.22331/q-2022-05-09-709

[6] T. Häner, D. S. Steiger, M. Smelyanskiy, and M. Troyer, "High performance emulation of quantum circuits," in *SC '16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2016, pp. 866–874.

[7] L. Dagum and R. Menon, "Openmp: An industry-standard api for shared-memory programming," *IEEE Comput. Sci. Eng.*, vol. 5, no. 1, pp. 46–55, Jan. 1998. [Online]. Available: https://doi.org/10.1109/99.660313

[8] C. R. Trott, D. Lebrun-Grandié, D. Arndt, J. Ciesko, V. Dang, N. Ellingwood, R. Gayatri, E. Harvey, D. S. Hollman, D. Ibanez, N. Liber, J. Madsen, J. Miles, D. Poliakoff, A. Powell, S. Rajamanickam, M. Simberg, D. Sunderland, B. Turcksin, and J. Wilke, "Kokkos 3: Programming model extensions for the exascale era," *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 4, pp. 805–817, 2022.

[9] D. Beckingsale, R. Hornung, T. Scogland, and A. Vargas, "Performance portable c++ programming with raja," in *Proceedings of the 24th Symposium on Principles and Practice of Parallel Programming*, ser. PPoPP '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 455–456. [Online]. Available: https://doi.org/10.1145/3293883.3302577

[10] A. Hayashi, A. Adams, J. Young, A. McCaskey, E. Dumitrescu, V. Sarkar, and T. M. Conte, "Enabling multi-threading in heterogeneous quantum-classical programming models," 2023. [Online]. Available: https://doi.org/10.48550/arXiv.2301.11559

Evolving Quantum Programming Abstractions

Austin Adams
School of Computer Science
Georgia Institute of Technology
aja@gatech.edu

Jeffrey Young
School of Computer Science
Georgia Institute of Technology
jyoung9@gatech.edu

Thomas M. Conte
School of Computer Science
Georgia Institute of Technology
conte@gatech.edu

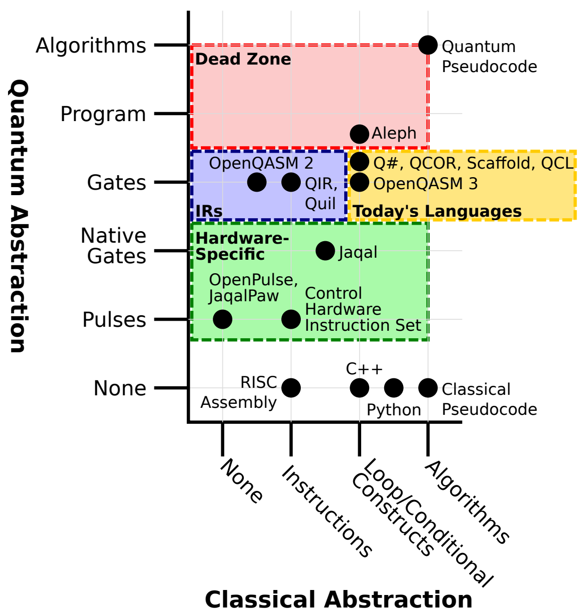

Fig. 1. Common quantum programming languages and IRs plotted by both their quantum and classical expressiveness. Qiskit [1] is notably missing because it is a library, not a programming language; however, it produces OpenQASM2 [2], an IR plotted above.

Topic: Quantum algorithms, applications, and compilation

I. Challenge

Quantum computers promise to find solutions to important problems exponentially faster, drawing their power from quantum mechanics. We argue that current quantum programming languages expose a level of abstraction largely unchanged since the 1980s, holding back the full capabilities of programmers, compilers, and hardware. Specifically, **realizing a productive quantum-classical environment will require significant leaps in quantum abstractions and programming from the current state-of-the-art**.

Current quantum circuit models provide useful theoretical frameworks for small-scale evaluations but driving real quantum hardware requires a programming language that is still very close to classical hardware description languages (HDLs). With a few exceptions noted below, quantum programming languages thinly wrap the circuit programming model with classical constructs for easier circuit generation.

To more precisely characterize the abstraction level of a quantum programming language, we must consider both its quantum and classical expressiveness, since many quantum programming languages support both quantum and classical code. Fig. 1 plots common quantum programming languages along these two dimensions and groups them into four categories that we discuss in the following sections.

A. Hardware-Specific Languages

Hardware-specific languages grant researchers low-level access to hardware for experimentation but are often not intended to be written by hand. For example, Jaqal (Just Another Quantum Assembly Language) is a gate-level language specifically for ion trap testbeds, where each native gate is an instruction [3], [4]. It includes high-level classical loop constructs in the IR itself for building repetitive circuits; thus, it is feasible to write Jaqal by hand, although not required. JaqalPaw (Jaqal Pulses and Waveforms) is similar but instead targets modulated lasers for an ion trap testbed at Sandia National Laboratories [4], [5].

B. Intermediate Representations

Hardware-agnostic intermediate representations (IRs) improve portability and overall convenience by expressing quantum circuits composed of general, non-native gates. Indeed, OpenQASM 2 [2] is perhaps the most common quantum IR due to its connection with the widely used Qiskit SDK. Its simple structure consists of a list of common gates and measurements, but it has limited support for classical logic, only supporting conditionally executing a quantum gate based on a previous measurement. Quil [6] and QIR [7] address the lack of support for classical logic in OpenQASM 2 by including some classical instructions. In particular, QIR supports all classical instructions in LLVM IR (a popular classical IR), leaving quantum operations as opaque function calls with quantum side effects.

C. Today's "High-Level" Languages

Purportedly higher-level languages promise to simplify quantum programming with higher-level constructs, but they still have limitations, especially for hybrid quantum-classical systems. QCL is the original quantum programming language [8], and its custom syntax allows programmers to write classical code that generates quantum circuits. Thus, QCL does not allow programmers to express classical code that runs during the qubit lifetime. Scaffold is similar except that the classical language is C++, which is much more familiar to programmers [9]. QCOR is also a quantum–classical language also hosted in C++ [10], [11], but rather than outputting a quantum circuit, it is hooked into the XACC compilation framework for easier circuit execution [12].

Neither QCL, Scaffold, nor QCOR can express classical computation during the qubit lifetime, since they focus on a "flat-circuit" model, but both Q# [13] and OpenQASM 3 [14] can. OpenQASM 3 is effectively a version of OpenQASM 2 with basic classical constructs such as arithmetic expressions, conditionals (e.g., `if`-statements), and loops (e.g., `for`-loops). In effect, OpenQASM 3 juxtaposes higher-level classical logic with low-level quantum gates. Q# attempts to more evenly match its level of classical abstraction with its level of quantum abstraction with a large library of helper functions for generating quantum gates.

D. Highest-Level Abstractions: The Dead Zone

We call the region of quantum abstraction above gates in Fig. 1 "the dead zone" since few programming languages even attempt to address the problem of quantum programming without requiring at least some knowledge of quantum mechanics. While still in-progress, Aleph [15] makes a notable contribution by requiring no physics knowledge from programmers. However, rather than exposing primitives for high-level algorithmic design, Aleph exposes Grover's algorithm itself as a built-in subroutine. Consequently, Aleph functions more as a tool for synthesizing circuits implementing Grover's algorithm than as a general-purpose quantum programming language.

II. Opportunities

Work on high-level quantum-classical languages and tools like QCOR/XACC, OpenQASM3, and Q# provide a good basis for future work in raising the abstraction layer of quantum programming by supporting IR translations from a common high-level IR to different quantum hardware backends. However, there is additional work that can be done to further improve the programmability of future hybrid systems.

Similar to how recent high-level synthesis tools for Field Programmable Gate Arrays like OneAPI and Vitis HLS allow programmers to use GPU-like syntax to design and implement low-level hardware implementations, we believe there is likely a viable abstraction for a quantum Domain Specific Language (DSL) that would be more approachable to the vast majority of computer programmers and early-career physicists. This is a codesign effort that requires intense effort in not just designing higher-level abstractions but also in compiler techniques to effectively lower these abstractions to gate-, pulse- and control hardware-level quantum backends.

A stronger quantum programming language would allow programmers to express not only the *how* of their programs through low-level gates, but also the *why*. This *why* could be some combination of specifying the higher-level linear algebra ultimately expressed by quantum algorithms and applying algorithmic steps (e.g., phase kickback) from a set of predefined library, similar to how most classical languages define arithmetic or string operations.

III. Assessment

Potential success would be demonstrated by implementing common quantum algorithms and small applications in a fashion that requires minimal expression of the algorithm with quantum-level gates. We estimate that these effects could be measured with respect to current high-level languages including OpenQASM3 and would be demonstrated with key algorithms like Grover's, Shor's, as well as with common scientific applications like Variational Quantum Eigensolvers (VQEs) and Computational Fluid Dynamics (CFDs).

IV. Timeliness and Maturity

The timeliness of this effort is boosted by DoE's recent investments in the quantum software and compiler space led by efforts like QCOR [10], [12] and Jaqal [5]. Current LLVM IR–based tools provide the needed substrate on which to design and implement an effective high-level quantum DSL, which could have wide-ranging impacts on the further adoption of quantum computing for scientific applications.

References

[1] Qiskit contributors, *Qiskit: An open-source framework for quantum computing*, 2023. DOI: 10.5281/zenodo.2573505.

[2] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, "Open quantum assembly language," *arXiv:1707.03429 [quant-ph]*, Jul. 13, 2017. arXiv: 1707.03429.

[3] B. C. A. Morrison, A. J. Landahl, D. S. Lobser, *et al.*, "Just another quantum assembly language (jaqal)," in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Oct. 2020, pp. 402–408.

[4] S. M. Clark, D. Lobser, M. C. Revelle, *et al.*, "Engineering the quantum scientific computing open user testbed," *IEEE Transactions on Quantum Engineering*, vol. 2, pp. 1–32, 2021, ISSN: 2689-1808.

[5] D. Lobser, J. Goldberg, A. J. Landahl, *et al.*, "JaqalPaw: A guide to defining pulses and waveforms for jaqal," Sandia National Laboratories, Apr. 2021, p. 22.

[6] R. S. Smith, M. J. Curtis, and W. J. Zeng, "A practical quantum instruction set architecture," arXiv, arXiv:1608.03355, Feb. 17, 2017, type: article. arXiv: 1608.03355[quant-ph]. (visited on 09/17/2022).

[7] A. Geller, *Introducing quantum intermediate representation (QIR)*, Sep. 23, 2020. [Online]. Available: https://devblogs.microsoft.com/qsharp/introducing-quantum-intermediate-representation-qir/.

[8] B. Ömer, "Quantum programming in QCL," Ph.D. dissertation, Technical University of Vienna, 2000.

[9] A. J. Abhari, A. Faruque, M. J. Dousti, *et al.*, "Scaffold: Quantum programming language (tech report)."

[10] T. M. Mintz, A. J. McCaskey, E. F. Dumitrescu, S. V. Moore, S. Powers, and P. Lougovski, "QCOR: A language extension specification for the heterogeneous quantum-classical model of computation," *ACM Journal on Emerging Technologies in Computing Systems*, vol. 16, no. 2, 22:1–22:17, Mar. 18, 2020, ISSN: 1550-4832.

[11] A. Mccaskey, T. Nguyen, A. Santana, D. Claudino, T. Kharazi, and H. Finkel, "Extending c++ for heterogeneous quantum-classical computing," *ACM Transactions on Quantum Computing*, vol. 2, no. 2, 6:1–6:36, Jul. 9, 2021, ISSN: 2643-6809.

[12] A. J. McCaskey, D. I. Lyakh, E. F. Dumitrescu, S. S. Powers, and T. S. Humble, "XACC: A system-level software infrastructure for heterogeneous quantum–classical computing," *Quantum Science and Technology*, vol. 5, no. 2, p. 024 002, Feb. 2020, ISSN: 2058-9565. DOI: 10.1088/2058-9565/ab6bf6. (visited on 02/03/2021).

[13] K. Svore, A. Geller, M. Troyer, *et al.*, "Q#: Enabling scalable quantum computing and development with a high-level DSL," ser. RWDSL2018, New York, NY, USA, Feb. 24, 2018, pp. 1–10.

[14] A. Cross, A. Javadi-Abhari, T. Alexander, *et al.*, "OpenQASM 3: A broader and deeper quantum assembly language," *ACM Transactions on Quantum Computing*, vol. 3, no. 3, 12:1–12:50, Sep. 6, 2022, ISSN: 2643-6809.

[15] A. Paz, *Aleph*. [Online]. Available: https://github.com/anpaz/aleph.

Toward Packet-switched Entanglement Distribution Networks

Ruozhou Yu (ryu5@ncsu.edu), Assistant Professor, *North Carolina State University*

Topic. Under the "models" emphasis area, this position paper introduces a new ***quantum network model and architecture*** for entanglement distribution networks (EDNs). The EDN architecture is built atop a novel *packet switching* paradigm, drastically different from existing *circuit switching*-based quantum networking and EDN architectures. We will argue that it brings significant benefits in terms of scalability, fault tolerance, throughput and efficient use of quantum resources, as well as many new challenges.

The central thesis of this paper is to argue that **packet switching** is likely to supersede circuit switching and become a mainstream architectural paradigm for EDNs with near-term quantum devices. This opens doors for exciting new opportunities but also substantial challenges prompting research and development.

Motivation. Distribution of entangled states (such as Bell states) is central to scalable quantum computing, communication and sensing systems. The current practice follows an early circuit switching paradigm, assigning fixed, dedicated resources (quantum channels along entanglement paths) to source-destination pairs. Its inefficacy comes from the lack of statistical multiplexing, inability for fault tolerance, and overhead for circuit (connection) management limiting scalability. While these generally hold true only in *large-scale* classical networks, the first two limitations exist and have a substantial impact even in a *small-scale* EDN with a few users and a few links. The core reason lies in the intrinsic *probabilistic nature* of quantum network operations, *e.g.*, the probabilistic generation process of spontaneous parametric down-conversion (SPDC) and probabilistic Bell state measurement (BSM) with linear optics.

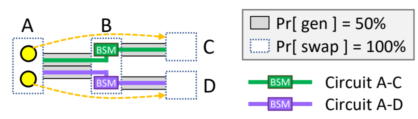

Figure 1: Two-hop EDN example.

Consider a small-scale EDN in Fig. 1 with two end-to-end tasks $A \to C$ and $A \to D$ via repeater B. With circuit switching, each task occupies one channel on link A-B and a downstream channel on B-C or B-D. The expected end-to-end entanglement rate for both circuits is $2 \times 0.5^2 = 0.5$ (both channels must succeed in each circuit). With statistical multiplexing, tasks can share entanglements successfully created on either channel along A-B, increasing expectation to 0.625 (25% increase). If temporary storage of generated entanglements (*i.e.*, quantum buffers) are available, the long-term expectation can reach 1.0, marking an 100% improvement. This simple example demonstrates the crucial advantage of both statistical multiplexing and buffering in building scalable EDNs.

Challenge. First, a new protocol stack needs to be designed to enable packet switching at various levels. Roughly, a packet contains a single qubit that belongs to a bipartite or multipartite entangled state, plus classical information for identification (entanglement ID, entanglement type, qubit index, etc.) and switching (*e.g.*, swapping/routing). At physical layer, one packet contains one entangled photon as its quantum payload, plus a classical wrapper (header and trailer) encoded as optical signal [1]. *How to properly multiplex and demultiplex classical and quantum information on an optical link to avoid interference?* The link layer switching requires processing classical header while holding the quantum payload in storage. *How fast must classical processing be to avoid decoherence of the quantum payload and how to achieve it? And how well can near-term quantum storage technologies hold the payload to wait for classical processing?* The link layer distribution is limited to a few kilometers in ground-based transmission, and will have to rely on either air/space-based links (such as satellite links) or the network layer to achieve long-distance distribution.

Up from the network layer is where things are going much differently from classical networks. Because of the no-cloning theorem, actual quantum operations such as quantum teleportation (or entanglement swapping) must be applied to form entanglements over long distances. On a high level, this means **consuming** link layer packets to generate long-distance entanglements. Consider four packets belonging to two pairs of Bell states. A swapping consumes two of the four packets, and upon success, forms an entanglement between the remaining two packets that were previously not entangled. Interestingly, network layer or up does not involve any transmission of quantum information carriers (such as photons), but is based on merely local operation and classical communication (LOCC). This means that *classical communication likely becomes the bottleneck that must be addressed to ensure in-time coordination before entanglements decohere*. Forming a bipartite end-to-end entanglement applies swapping repeatedly along an "entanglement path", leading to the basic networking question: *how to route entanglement generation along optimal paths to maximize efficiency and quality of entanglement distribution?* The answer commonly resides in a complex space given various attributes of quantum links and nodes, such as probability of success, noise and error rate, node buffers, etc. A more general question involving *multipartite entanglement* would turn path finding into more general graph problems for which little is known about either complexity or algorithmic solutions. Further, packet switching offers the opportunity of freely sharing entanglements among connections (statistical multiplexing) and buffering, leading to many new questions such as *1) how to enable distributed control and dynamic routing of entanglement flows*, *2) how to ensure high quality-of-service (QoS) such as high fidelity of the distributed entanglements*, *3) how to efficiently and intelligently manage quantum buffers at each intermediate node*, and more. The fundamental buffering issue further leads to the question of congestion control that classically resides in the transport layer: *how to dynamically adjust network-wide entanglement generation and swapping behaviors to avoid excessive queueing (and hence decoherence) at repeater nodes*. In the end, all above questions merge into one question: *how these pieces can be designed to support end-to-end entanglement distribution that can satisfy the needs of various quantum applications* — a question that must be answered with those applications in mind!

Beyond the stack itself, progress is needed in several critical enabler technologies to implement a packet-switched EDN. *Quantum memories* (though classically addressed) are essential for both switching and buffering, and their size, success probability, storage time and storage fidelity have substantial impact on network performance. Given inevitable noise and imperfect quantum devices, *quantum error correction* (including *entanglement distillation*) must be conducted at all layers, even including the network layer where error control is not concerned with in classical networks. The *optimal EDN architecture* depends on availability and capacity of devices, efficiency of processes, and noise and error rate of the network environment. When the goal is to simultaneously support multiple types of entanglement (bipartite, multipartite, entangled qudits or continuous-variable) with one EDN, the *co-existence, multiplexing/demultiplexing, storage/buffering, scheduling, and transduction of multi-type entanglements* pose a much larger class of questions to answer, especially when any of these can happen during any stage of packet switching on-the-fly. Considering near-term quantum devices, *what is the achievable advantage of a packet-switched EDN compared to a circuit-switched one?*

Opportunity. The devices (such as quantum memories and repeaters) themselves need both physics and engineering efforts. On top of that, new advances are required in: **mathematical models** for devices, processes across the stack, noise, network-wide operations, architectural characterization, and application requirements and demands; **protocols and algorithms** for efficiently and optimally controlling and coordinating operations across the stack and among nodes; **simulation and testing** environments and benchmarks to evaluate the efficacy of architectural and protocol designs. Each of these require a combination of highly interdisciplinary expertise from several domains and across the stack: physicists (theoretical or experimental), mathematicians, theoretical computer scientists, network and system researchers, and network engineers.

Assessment. Assessment of quantum network/EDN architecture itself is an open and important problem. Expectedly, theoretical analysis and simulation-based study will be primary tools for assessment in early stage, while emulation and testbed demonstration will become valuable when corresponding technologies mature. While evaluation metrics should closely tie to the applications that the EDN supports, common metrics including the entanglement distribution rate (EDR), the quality (fidelity) of entanglements, and resource costs (such as buffer occupancy) are valuable indications of the success of an EDN architecture. It is expected (and preliminarily validated in our study [2, 3]) that a packet-switched EDN can achieve both higher EDR and fidelity than circuit switching. More comprehensive study and development are needed.

Timeliness. The proliferation of quantum networking/EDN research was triggered by recent success in building lab-scale quantum network testbeds. Prototypes of long-term optical quantum memories and quantum repeaters are expected to emerge in a few years [4]. Much like in the initial decade of Internet (ARPANET)'s invention, the study of scalable and efficient EDN architecture will set forth the foundation of the real deployment and expansion of EDN into a fully fledged Quantum Internet. Any decision made by vendors/standardization groups/the community at this stage will likely have a profound impact on what it can eventually achieve and enable.

[1] S. Diadamo, B. Qi, G. Miller, R. Kompella, and A. Shabani, "Packet Switching in Quantum Networks: A Path to the Quantum Internet," *Physical Review Research*, vol. 4, no. 4, pp. 1–17, 2022.

[2] H. Gu, Z. Li, R. Yu, X. Wang, F. Zhou, and J. Liu, "FENDI: High-Fidelity Entanglement Distribution in the Quantum Internet," *arXiv:2301.08269*, pp. 1–13, 2023.

[3] H. Gu, R. Yu, Z. Li, X. Wang, and F. Zhou, "ESDI: Entanglement Scheduling and Distribution in the Quantum Internet," *arXiv:2303.17540*, pp. 1–10, 2023.

[4] Y. Wang, A. N. Craddock, R. Sekelsky, M. Flament, and M. Namazi, "Field-Deployable Quantum Memory for Quantum Networking," *Physical Review Applied*, vol. 18, no. 4, p. 044058, oct 2022.

Quantum Algorithm, Networking, and Sensing Technologies for Future Power Grids

M. Yue[a], S. Yoo[a], Y. Chen[b], T. Zhao[a], and J. Wang[c]

[a]Brookhaven National Laboratory, [b]Pacific Northwest National Laboratory, [c]Southern Methodist University

Emails: {yuemeng,sjyoo,tzhao1}@bnl.gov,Yousu.Chen@pnnl.gov,jianhui@mail.smu.edu

Topic: applications, models, and algorithms

Challenge: The increasing renewable generation introduces faster dynamics of the microsecond timescale of power electronics devices interfacing the renewables with the grids. While traditional supervisory Control and Data Acquisition (SCADA, usually taking samples every 2 to 4 seconds) and Phasor Measurement Unit (PMU, ~30 samples per second) types of measurements have been available and applied for grid operation, for the grid operators and planners, having access to higher resolution and time-synchronized measurements is not only beneficial but also can be necessary to better understand the dynamic behaviors of the power grid. For example, the emerging point-on-wave (POW) data technology (e.g., 256 samples/second or higher) [1] can reveal more about local and wide-area conditions. However, only waveforms are the true representation of the dynamic behaviors of systems and components [2]. In addition, the electric power grid stretches across several thousands of miles, and the inherent data latency has always been an issue in real-time applications. Given the grand challenges for better understanding the fast dynamics and protection and control, it is envisioned that the existing sensing technologies are insufficient for high-speed, time-synchronized, and wide-area applications. Next-generation sensing and networking technologies are needed for the future power grid.

Opportunity: Quantum networking (QN) can facilitate the transmission of information in the form of qubits between different quantum processors and/or types of quantum systems at different locations. QN technologies will not only enable distributed computation and quantum sensing (QS) but also enhance communication speed, safety, and security. Notably, emerging QS [3] can achieve granularity beyond the limit of classical sensing technologies by defining the measurement of physical quantities based on quantum objects, quantum coherence, or quantum entanglement [4]. QS optimally estimates classical parameters encoded in quantum transformations, offering unprecedented combinations of range, resolution, and sensitivity for measuring parameters, enhancing the sensitivity, and providing the most authentic and granular measurements for the parameters of interest in the physical process. A combination of scalable quantum fiber-optic networking and distributed QS and algorithms can achieve superior performance for temporal and spatial scales ranging, i.e., from seconds to femtoseconds (~10^{-15} seconds) and from thousands of kilometers to micrometers, offering a promising solution to addressing the challenges in the power grid.

There are many needs and opportunities to use high-resolution, time-synchronized power system data. Some potential applications are discussed here:

- The increasing penetration level of renewables and inverter-based resources (IBRs) poses significant threats to grid stability. To better understand the dynamic behaviors of the power grid and the coordination of protection and control systems of IBRs, high temporal resolution measurements and waveforms across a large area will be needed due to the fast IBR dynamics and the distributed nature of renewables. Differential protection schemes can be a highly effective solution to addressing the protection issue under high IBRs by comparing the currents at both ends of a transmission line. Its implementation requires accurate measurements of currents.

- Sub-synchronous resonance (SSR) is coincident oscillation at a natural harmonic frequency lower than the system's nominal frequency (60 Hz) that can lead to devastating impacts, harming resonating transmission elements, fracturing a generator shaft, and causing cascading outages. SSR oscillations can only be observed in high-speed waveforms. Harmonics and power quality analysis will be enabled on a much larger scale with high-speed measuring devices.
- A geomagnetic disturbance may produce geomagnetically-induced currents (GICs) that create high-frequency harmonics and voltage differentials at different electric transmission ground points, leading to thermal stress and damage to grid components. High-order harmonics and high-altitude electromagnetic pulse (HEMP) cannot be captured and detected by the existing PMU technologies. This will be enabled and/or facilitated by using electric and magnetic field quantum sensors such as SQUID (Superconducting Quantum Interference Device) magnetometers and atomic magnetometers [5].

There can be many more potential applications that are only possible by the QN and distributed QS technologies and need to be identified. In addition, quantum algorithm-based methods and tools need to be developed for implementing specific grid applications based on the QS data. This also calls for innovative solutions in data storage, signal processing, and data analytics.

Assessment: The success of QN, algorithms and distributed QS applications can be measured in terms of whether specific grid application goals can be affordably achieved only with the QN and distributed QS technologies and the performance improvement (speed/latency, accuracy, scalability, and security) compared to the state of practices.

Timeliness or maturity: The rapid transformation of the electric power grid will benefit from better situational awareness driven by measurement data. Significant progress has been achieved in QN and distributed QS technologies. For example, quantum sensors may be commercially available in 3 to 5 years [5]. In addition, experience from many active applications of QS and QN in areas such as X-ray microscopy, lidar, and telescope (astrometry) can be leveraged. The QN, algorithms, and QS technologies may fundamentally shift the paradigm for observability and controllability in the future grid.

References:

[1] A. Silverstein, J. Follum, "High-resolution, time-synchronized Grid Monitoring Devices," North American Synchrophasor Initiative Technical Report, March 20, 2020.

[2] W. Xu, et al, "Synchronized Waveforms – A Frontier of Data-Based Power System and Apparatus Monitoring, Protection, and Control," IEEE Transactions on Power Delivery, Vol. 37, No. 1, February 2022.

[3] M. Hayashi, Quantum Information Theory: Mathematical Foundation, Springer, Berlin 2017.

[4] S. E. Crawford, et al, "Quantum Sensing for Energy Applications: Review and Perspective," Adv. Quantum Technol. 2021, 4, 2100049.

[5] "Quantum Sensing Use Cases: Prospects and Priorities for Emerging Quantum Sensors," QED-C, September 2022.

Making quantum error mitigation practical

William J. Zeng[1,*] Nathan Shammah[1]
[1] *Unitary Fund*

Useful applications of quantum computers require significant reductions in logical error rates [1–6]. One direction to achieve this is to implement quantum error correcting codes. Complementing quantum error correction, are new techniques for quantum error mitigation [7–17]. These are algorithmic methods that are designed to be less experimentally demanding than full quantum error correction. However, this comes at the cost of being less general and more heuristic.

I. Challenges in quantum error mitigation

There are several key challenges in making error mitigation practical.

Reducing error mitigation overhead. For example, in some techniques, the number of samples N required to approximate the expectation value output from an ideal quantum computer to within an error δ scales [18] as $N \propto \gamma^2/\delta^2$, where γ is a constant that becomes larger as the quantum program becomes larger and the quantum computer becomes noisier. The γ values of approximately 1.02 have been measured in IBM processors [19]. This exponential dependence emphasizes how important it is to improve performance of different error mitigating techniques and to study their fundamental limits [20].

Calibrating optimal techniques. While there are a growing number of options available, this means the programmer must choose what techniques to use and with what parameters. Making this choice well depends on the hardware target and on having a good model of the noise. Further, there is a tradeoff between spending valuable quantum computer time further calibrating the error migitation vs. exploiting the model that is currently available. Additionally, while there have been shown benefits to composing error mitigating techniques—such as [13] where generalizing PEC and ZNE produces a more robust method—there are open research questions about how best to do this composition. These calibration and composition choices need to be made scalable so that they apply to larger QPUs whose output cannot be simulated and to problems where we cannot train on a previously known answer. Finally, several error mitigating techniques require lower level access to control electronics that is not always available from vendors. More abstract techniques and the integration of error mitigation at lower levels of the stack are needed to improve performance.

* will@unitary.fund

Error mitigation and fault-tolerance. How can error mitigation be applied to accelerate the deployment of error correcting codes? For example, Pauli twirling can convert coherent errors into stochastic noise [21] that could improve the performance of error correction. Further, error mitigation can be extended into the fault-tolerant regime where it can reduce overheads [22] and, in some examples, improve the number of logical operations that can be applied by a factor of 1000X [23].

II. Opportunities for quantum error mitigation

These challenges are opportunities to both improve the performance of today's quantum computers and also accelerate roadmaps across hardware modalities, including quantum sensors and networks. If properly seized, then error mitigation can provide a smooth ramp up towards quantum advantage [19], making it easier for the quantum technology industry to cross the chasm to valuable applications. We describe three key categories of opportunity:

There is an opportunity to use **open source software**, such as the cross platform error-mitigating compiler Mitiq [24], to study and automate the calibrations needed for optimal error mitigation. Open source error mitigation implementations are accretive, allowing researchers and programmers to take advantage of the state of the art without needing to implement everything from scratch themselves. The community using this software can study and fine tune these techniques across hardware platforms and upstream their learning.

Integrating these error mitigating techniques with hardware design offers an opportunity for **hardware-software co-design**. Here, error mitigating techniques can be considered in both NISQ and fault-tolerant quantum computer architectures. One could, for example, tailor the noise channels towards ones that are easy for mitigating techniques to calibrate and counter.

Research at the **intersection of error mitigation and error correction**. As error correction becomes more practical, it is likely that there are new error mitigating techniques that can be discovered that integrate well with error correction.

III. Assessment and Timeline

Progress on error mitigation can be assessed using benchmarks of performance such as effective quantum volume [25], improved performance of application level benchmarks, or improvements in logical gate fidelity or coherence. It is important for these assessments that performance takes into account the cost and time of classical post- and pre- computations used in the error mitigation. Ideally these assessments of mitigation performance will occur in the supremacy regime where it is non-trivial (or impossible) to classically simulate the results directly. A final assessment for software tools, such as error mitigating compilers, is their usage by the community with metrics like downloads, github stars, citations, etc.

Now is a good time to focus on these error mitigation challenges since (1) we have a stable pool of techniques that are ready to be reduced to practice and (2) we have a need from applications and fault-tolerant design to reduce error rates as quickly as possible. Success on these challenges can meaningfully affect the timeline to useful quantum computing across the whole field.

---

[1] Alexander M Dalzell, B David Clader, Grant Salton, Mario Berta, Cedric Yen-Yu Lin, David A Bader, Nikitas Stamatopoulos, Martin JA Schuetz, Fernando GSL Brandão, Helmut G Katzgraber, et al. End-to-end resource analysis for quantum interior point methods and portfolio optimization. *arXiv preprint arXiv:2211.12489*, 2022.

[2] Michael Kühn, Sebastian Zanker, Peter Deglmann, Michael Marthaler, and Horst Weiß. Accuracy and resource estimations for quantum chemistry on a near-term quantum computer. *J. Chem. Theor. Comp.*, 15(9):4764–4780, 2019.

[3] Shouvanik Chakrabarti, Rajiv Krishnakumar, Guglielmo Mazzola, Nikitas Stamatopoulos, Stefan Woerner, and William J Zeng. A threshold for quantum advantage in derivative pricing. *Quantum*, 5:463, 2021.

[4] Andrew J Daley, Immanuel Bloch, Christian Kokail, Stuart Flannigan, Natalie Pearson, Matthias Troyer, and Peter Zoller. Practical quantum advantage in quantum simulation. *Nature*, 607(7920):667–676, 2022.

[5] Michael E Beverland, Prakash Murali, Matthias Troyer, Krysta M Svore, Torsten Hoeffler, Vadym Kliuchnikov, Guang Hao Low, Mathias Soeken, Aarthi Sundaram, and Alexander Vaschillo. Assessing requirements to scale to practical quantum advantage. *arXiv preprint arXiv:2211.07629*, 2022.

[6] Craig Gidney and Martin Ekerå. How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits. *Quantum*, 5:433, 2021.

[7] Suguru Endo, Zhenyu Cai, Simon C Benjamin, and Xiao Yuan. Hybrid quantum-classical algorithms and quantum error mitigation. *J. Phys. Soc. Jap.*, 90(3):032001, 2021.

[8] Zhenyu Cai, Ryan Babbush, Simon C. Benjamin, Suguru Endo, William J. Huggins, Ying Li, Jarrod R. McClean, and Thomas E. O'Brien. Quantum error mitigation. *arXiv preprint arXiv:2210.00921*, 2022.

[9] Chao Song, Jing Cui, H Wang, J Hao, H Feng, and Ying Li. Quantum computation with universal error mitigation on a superconducting quantum processor. *Science advances*, 5(9):eaaw5686, 2019.

[10] Christophe Vuillot. Is error detection helpful on IBM 5Q chips? *arXiv preprint arXiv:1705.08957*, 2017.

[11] Abhinav Kandala, Kristan Temme, Antonio D Córcoles, Antonio Mezzacapo, Jerry M Chow, and Jay M Gambetta. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567(7749):491–495, 2019.

[12] Tudor Giurgica-Tiron, Yousef Hindy, Ryan LaRose, Andrea Mari, and William J Zeng. Digital zero noise extrapolation for quantum error mitigation. In *Int. Conf. Quant. Comp. Eng.*, pages 306–316. IEEE, 2020.

[13] Andrea Mari, Nathan Shammah, and William J Zeng. Extending quantum probabilistic error cancellation by noise scaling. *Physical Review A*, 104(5):052607, 2021.

[14] Google AI Quantum et al. Hartree-Fock on a superconducting qubit quantum computer. *Science*, 369(6507):1084–1089, 2020.

[15] Benjamin McDonough, Andrea Mari, Nathan Shammah, Nathaniel T Stemen, Misty Wahl, William J Zeng, and Peter P Orth. Automated quantum error mitigation based on probabilistic error reduction. *arXiv preprint arXiv:2210.08611*, 2022.

[16] Shuaining Zhang, Yao Lu, Kuan Zhang, Wentao Chen, Ying Li, Jing-Ning Zhang, and Kihwan Kim. Error-mitigated quantum gates exceeding physical fidelities in a trapped-ion system. *Nature Commun.*, 11(1):1–8, 2020.

[17] Vincent Russo, Andrea Mari, Nathan Shammah, Ryan LaRose, and William J Zeng. Testing platform-independent quantum error mitigation on noisy quantum computers. *arXiv preprint arXiv:2210.07194*, 2022.

[18] Ryuji Takagi. Optimal resource cost for error mitigation. *Physical Review Research*, 3(3):033178, 2021.

[19] IBMResearch. With fault tolerance the ultimate goal, error mitigation is the path that gets quantum computing to usefulness, 2021.

[20] Ryuji Takagi, Suguru Endo, Shintaro Minagawa, and Mile Gu. Fundamental limits of quantum error mitigation. *npj Quantum Information*, 8(1):114, 2022.

[21] Akel Hashim, Ravi K. Naik, Alexis Morvan, Jean-Loup Ville, Bradley Mitchell, John Mark Kreikebaum, Marc Davis, Ethan Smith, Costin Iancu, Kevin P. O'Brien, Ian Hincks, Joel J. Wallman, Joseph Emerson, and Irfan Siddiqi. Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor. *Phys. Rev. X*, 11:041039, Nov 2021.

[22] Matteo Lostaglio and Alessandro Ciani. Error mitigation and quantum-assisted simulation in the error corrected regime. *Phys. Rev. Lett.*, 127(20):200506, 2021.

[23] Yasunari Suzuki, Suguru Endo, Keisuke Fujii, and Yuuki Tokunaga. Quantum error mitigation as a universal error reduction technique: Applications from the NISQ to the fault-tolerant quantum computing eras. *PRX Quantum*, 3(1):010345, 2022.

[24] Ryan LaRose, Andrea Mari, Sarah Kaiser, Peter J Karalekas, Andre A Alves, Piotr Czarnik, Mohamed El Mandouh, Max H Gordon, Yousef Hindy, Aaron Robertson, et al. Mitiq: A software package for error mitigation on noisy quantum computers. *Quantum*, 6:774, 2022.

[25] Ryan LaRose, Andrea Mari, Vincent Russo, Dan Strano, and William J Zeng. Error mitigation increases the effective quantum volume of quantum computers. *arXiv preprint arXiv:2203.05489*, 2022.

Designing Quantum Routers for Quantum Internet

Inder Monga
ESnet, LBNL
imonga@es.net

Liang Zhang
ESnet, LBNL
lzhang9@es.net

Yufeng Xin
RENCI, University of North Carolina
yxin@renci.org

1 Topic: quantum networks, models, algorithms

2 Challenge

Enabling quantum communication between any two locations on Earth is the primary objective of the quantum internet. To achieve this, the role of quantum routers is critical. However, these routers are presently in their nascent stages and encounter several obstacles that need to be addressed in the forthcoming years.

2.1 Adaptability to heterogeneous quantum resources is crucial

Several qubit platforms shown in Figure 1(a), have been proposed for constructing quantum routers. Handling the varied hardware and software components poses a difficulty for quantum routers. Presently, numerous quantum protocols adopt a monolithic model that obstructs scalability and necessitates significant effort to accommodate the heterogeneous quantum resources that the router possesses. With network topologies becoming more expansive and intricate, limited resources like qubit memories and entanglements generated are prone to underutilization, which may lead to inadequate performance of quantum networks.

2.2 Accommodating multiple concurrent entanglement flows that require high rates and fidelity is vital

Quantum routers face a significant challenge in generating entanglement with high success rates and fidelity. As network size increases, numerous entanglement flows require quantum router service within a short coherence time window. The existing multiplexing schemes lacked the necessary guarantees on entanglement fidelity. Although switching-based schemes enhance fidelity, they can only accommodate quantum connections on the order of ten due to physical limitations, which is inadequate for the exponential growth of quantum networks.

2.3 Large-scale networks that feature highly dynamic quantum links require efficient entanglement routing algorithms

Quantum entanglement between neighboring nodes is probabilistic and has a fleeting lifespan. The resulting dynamic nature of quantum links renders routing algorithms based on global information or centralized control unsuitable. The OSPF routing protocol, which relies on the link states of the entire network, is a case in point and underperforms over quantum links. As a result, novel routing algorithms must be developed to operate effectively in large-scale quantum Internet.

3 Opportunity

3.1 Virtualization could facilitate the use of heterogeneous quantum resources

Virtualization is a powerful tool in classic networks, as it enables exclusive physical resources to be concealed and assigned to specific functions. Schoute [1] introduced the concept of virtual quantum links (VQL) to demonstrate shared entanglement. The VQL can be expanded into a Virtual Link layer, as shown in Figure 1(a), allowing classical algorithms and protocols to address emerging quantum challenges. Exploring the virtualization of quantum networks may lead to a network overlay that employs advanced resource allocation and optimization techniques.

3.2 Quantum router capacity could be expanded by constructing local switching network of small router modules

Lee's study [2] showcases the successful development of a photonic integrated circuit (PIC) that exhibits high fidelity and entanglement rates. To increase the practical capacity of quantum routers, one promising avenue of exploration is to create several tiers of interconnections between these chips, forming a local switching network (SwitchNet) in Figure 1(a). Popular network topologies such as Clos and Benes have already demonstrated desirable properties like non-blocking and cost-effectiveness when establishing connectivity for large-scale switches in traditional communication, and could be further developed to build high-capacity quantum routers in Figure 1(b)

3.3 Exploring novel algorithms could advance entanglement routing

Before being implemented on the quantum Internet, routing algorithms originally designed for classical networks have been re-evaluated. Researchers, such as Pant [3], have explored routing schemes that utilize the diversity of multiple paths, which have shown to outperform single path one in entanglement generating rates. Additionally, Shi [4] has investigated a group of distributed routing algorithms that solely rely on link information from nearby neighbors in Figure 1(c). His findings suggest that such algorithms exhibit significant improvements in performance compared to classical algorithms that depend on the link state of the entire network. Furthermore, optimization such as linear programming can also be employed to determine the optimal path for a given quantum connection constraint.

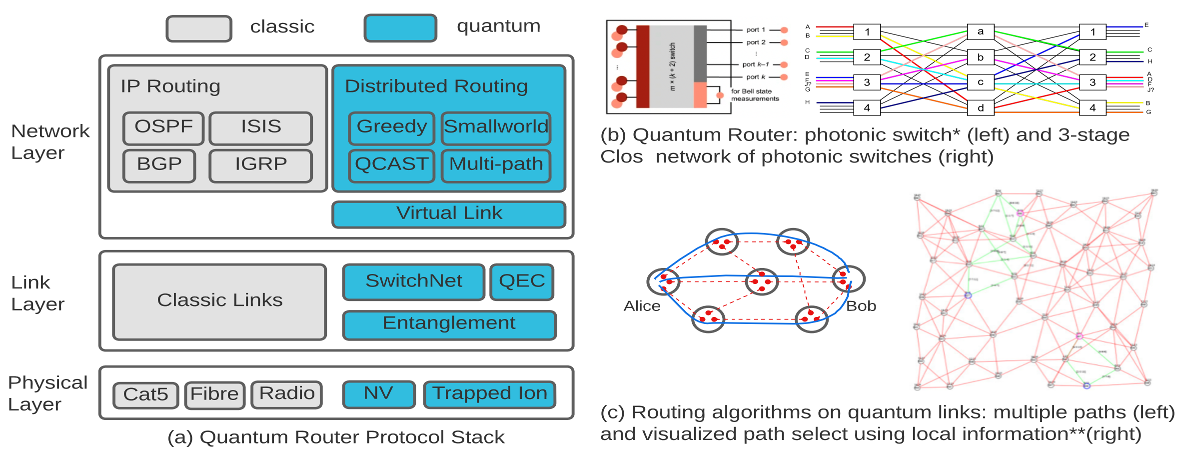

(a) Quantum Router Protocol Stack

(b) Quantum Router: photonic switch* (left) and 3-stage Clos network of photonic switches (right)

(c) Routing algorithms on quantum links: multiple paths (left) and visualized path select using local information**(right)

Figure 1: Quantum Router Design

4 Assessment

Quantum routers will undergo assessment of their functionality and performance using simulations and Quantum testbeds. The tests will encompass different network topologies, sizes, and noise levels, with crucial metrics being the end-to-end entanglement rates, fidelity, and capacity.

5 Timeliness or maturity

Currently, state-of-the-art quantum computers are equipped with processors containing over 400 qubits. However, quantum networks are predominantly designed for short distances and support only a restricted number of quantum nodes. Therefore, there is an urgent need to investigate quantum routers that can facilitate large-scale quantum communication.

References

[1] E. Schoute, L. Mancinska, T. Islam, I. Kerenidis, and S. Wehner, “Shortcuts to quantum network routing,” *CoRR*, 2016.

[2] Y. Lee, E. Bersin, A. Dahlberg, S. Wehner, and D. Englund, “A quantum router architecture for high-fidelity entanglement flows in quantum networks,” *npj Quantum Information*, 2022.

[3] M. Pant, H. Krovi, D. Towsley, L. Tassiulas, L. Jiang, P. Basu, D. Englund, and S. Guha, “Routing entanglement in the quantum internet,” *npj Quantum Information*, 2019.

[4] S. Shi and C. Qian, “Concurrent entanglement routing for quantum networks: Model and designs,” *Sigcomm*, 2020.

[1] [2]

Q-NAS: Quantum Noise-Resilience Ansatz Searching for Variational Quantum Algorithm

Junyao Zhang[1], Potok Thomas[2], Date, Prasanna[2], Yiran Chen[1]
[1]Duke University, [2]Oak Ridge National Laboratory
jz420@duke.com

I. Topic

We propose Q-NAS, a software algorithm that effectively identifies optimal, noise-resilient quantum circuits (ansatz), which shows promise in advancing variational quantum algorithms (VQA) [1] towards real-world applicability in the noise intermediate scale quantum (NISQ) [2] era.

II. Challenge

Quantum computing leverages quantum mechanical phenomena to provide computing advantages, thereby paving the way for groundbreaking advancements across various domains, such as chemistry [3], fundamental software algorithms [4], and machine learning [5]. VQA, a promising NISQ algorithm, has been pivotal in resolving classically intractable problems in the aforementioned domains with its various applications, such as Variational Quantum Eigensolver (VQE) for chemistry, Quantum Approximate Optimization Algorithm (QAOA) for approximation, and quantum machine learning for machine learning. VQA is a long-running iterative algorithm that deploys a classical optimizer to train a parameterized quantum circuit on a quantum machine. The circuit parameters are tuned in each iteration to approach the targets of application (usually minimization problems), such as estimating the ground state energy of molecules.

However, the presence of noise in quantum systems hinders the practical implementation of VQA [2]. Qubits are susceptible to noise from various sources, including environmental factors [6], manufacture device defects [7] and interactions with other quantum systems [7], and imperfections in quantum hardware and control systems [8]. Noise errors from these sources compound and accumulate, thus increasing the probability of erroneous outcomes in quantum circuits (ansatz). Fig.1 illustrates the impact of noise over a circuit batch (25 identical circuits) from IBMQ Belem, where the average fidelity variation over 100 executions reaches 42.8%. The cumulative effect of noise errors on quantum circuits further compromises the reliability of VQAs. A classification application using measurements on IBMQ Lagos and noise-free simulation in Fig.2 highlights two key observations: (1) More parameters improve the ability of the model but also introduce more noise, offsetting benefits (accuracy peak in a quantum device is at 44 parameters), and (2) quantum noise exacerbates performance variance. These findings emphasize the necessity for noise-resilience ansatz search in designing robust circuits.

III. Opportunity

To facilitate higher-level optimization, our goal is to identify optimal noise-resilience quantum circuits and their corresponding qubit mappings, thereby improving the performance of targeted tasks on specific quantum devices. This optimization poses a significant challenge to algorithmic scalability, since solving a two-level optimization problem is computationally intensive that requires iterative circuit sampling, parameter training, and evaluation within a vast design space. To tackle this problem, we introduce Q-NAS, a novel predictor-based approach for co-searching ansatz and qubit mapping, as depicted in Fig.3.

This approach entails training a performance predictor on a limited number of samples. The predictor subsequently acts as a surrogate model for ground-truth performance, guiding the ansatz search throughout the entire design space. The entire framework comprises the following steps: (a) To create a comprehensive dataset covering an extensive search space, we stack layers of various pre-defined parameterized gates to design ansatz and pair ansatz with different qubit mappings; (b) We

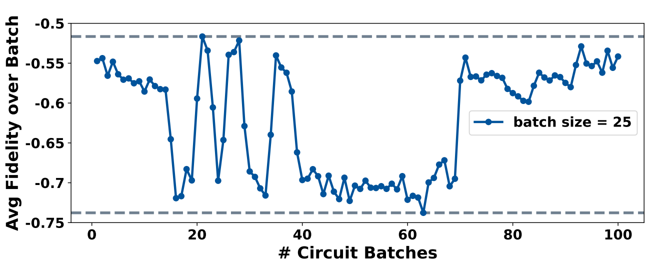

Fig. 1. Noise errors on circuits (ansatz). Circuit data are collected by 100 continuous runs of a circuit batch from an experiment run on IBMQ Belem. Each data point is the average expectation values of the circuit batch (25 identical circuits).

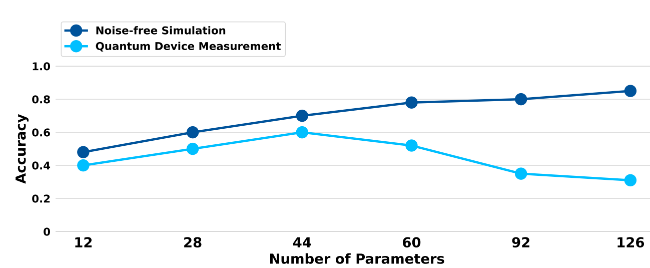

Fig. 2. Noise errors on classification application (run on noise free simulator and IBMQ Lagos device), More parameters increase the noise-free accuracy but degrade measured accuracy due to larger gate errors.

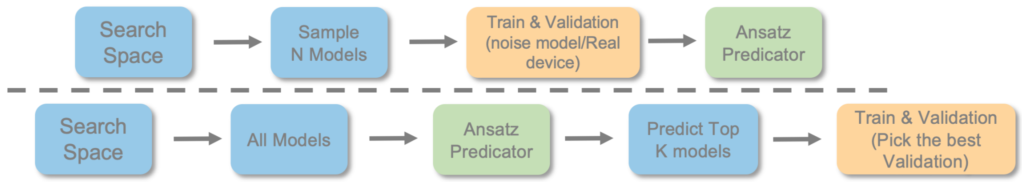

Fig. 3. Noise-resilience ansatz searching. A predictor is first trained from sampled models. Then, evolutionary search is performed to find the most robust ansatz.

randomly sample a set of models from the dataset, optimize their parameters, and assess their performance using noise-aware simulators or real quantum hardware; (c) We employ a traditional loss function as the objective for predictor training on the sampled models; (d) Utilizing the trained predictor, we apply regularized evolution to effectively probe the entire search space and identify the top $\mathcal{M}$ models. We further select the best model based on its actual validation performance. We effectively decouple the training and searching processes with this approach, which incurs the search cost only associated with training a small dataset for the predictor. The predictor can then efficiently and reliably evaluate the entire dataset. Consequently, this method substantially reduces the search cost while preserving high accuracy in the evaluation process.

IV. Assessment

The primary evaluation metric for the framework is the ability of its selected ansatz to deliver superior performance compared to other ansatz's when deployed on real-world quantum devices. This entails achieving higher accuracy and robustness in the presence of noise. Another important criterion for evaluating the proposed framework is the associated cost of searching for the optimal ansatz. This includes the number of circuits trained and the number of iterations required for each training process. By minimizing these factors, the solution would demonstrate not only its effectiveness in terms of performance but also its efficiency in terms of computational resources and time. In summary, the proposed solution will be assessed based on its performance on real-world quantum devices and the cost-efficiency in finding the optimal ansatz. The combination of these metrics will provide a comprehensive evaluation of the solution's effectiveness and practical applicability.

V. Timeliness

Various technologies are proposed to actively address noise challenges in quantum systems, including quantum error-correcting codes [9], noise-adaptive quantum program compilation [10], qubit mapping [11], instruction scheduling for crosstalk mitigation [7]. Despite these advances, prior research has predominantly focused on the gate-level compilation for mitigating the noise impact, leaving higher-level optimization for noise-resilient quantum circuits relatively unexplored. Our work aims to address this gap by co-designing ansatz search and qubit mapping strategies to enhance the noise resilience of ansatz.

VI. Conclusion

In conclusion, VQA is poised to become a massive enabler for harnessing quantum power and driving significant advancements. However, noise in quantum systems obstructs the practical deployment of quantum algorithms and the demonstration of quantum computational supremacy. Additionally, current manufacturing limitations restrict our ability to address noise at the hardware level. This restriction emphasizes the importance of designing software algorithms to mitigate noise effectively. Our work shows promise in pushing NISQ era quantum machine learning toward real-world applicability.

References

[1] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, "Quantum optimization using variational algorithms on near-term quantum devices," *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, 2018.

[2] J. Preskill, "Quantum computing in the nisq era and beyond," *Quantum*, vol. 2, p. 79, 2018.

[3] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'brien, "A variational eigenvalue solver on a photonic quantum processor," *Nature communications*, vol. 5, no. 1, p. 4213, 2014.

[4] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.

[5] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.

[6] P. Murali, N. M. Linke, M. Martonosi, A. J. Abhari, N. H. Nguyen, and C. H. Alderete, "Full-stack, real-system quantum computer studies: Architectural comparisons and design insights," in *Proceedings of the 46th International Symposium on Computer Architecture*, 2019, pp. 527–540.

[7] P. Murali, D. C. McKay, M. Martonosi, and A. Javadi-Abhari, "Software mitigation of crosstalk on noisy intermediate-scale quantum computers," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 1001–1016.

[8] M. Sun and M. R. Geller, "Efficient characterization of correlated spam errors," *arXiv preprint arXiv:1810.10523*, 2018.

[9] D. A. Lidar and T. A. Brun, *Quantum error correction*. Cambridge university press, 2013.

[10] M. Alam, A. Ash-Saki, and S. Ghosh, "Circuit compilation methodologies for quantum approximate optimization algorithm," in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020, pp. 215–228.

[11] L. Liu and X. Dou, "Qucloud: A new qubit mapping mechanism for multi-programming quantum computing in cloud environment," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2021, pp. 167–178.