

Advanced Scientific Computing Research

Quantum Information Science Awards Abstracts

Quantum Performance Assessment

Robin Blume-Kohout, Erik Nielsen, Tim Proctor, Ken Rudinger, Kevin Young, Sandia National Laboratories

Efficient and Reliable Mapping of Quantum Computations onto Realistic Architectures

Andrew Childs, Alexey Gorshkov, Michael Hicks, Xiaodi Wu, Robert Rand, University of Maryland

Quantum Algorithms, Mathematics and Compilation Tools for Chemical Sciences

Wibe Albert de Jong, Lawrence Berkeley National Laboratory, Stefan Wild, Argonne National Laboratory, Alán Aspuru-Guzik, University of Toronto

Superconducting Quantum Hardware Architecture Pathfinder

Jonathan L DuBois, Eric T. Holland, Yaniv J. Rosen, Luis A. Martinez, Xian Wu
Quantum Coherent Device Physics Group, Lawrence Livermore National Lab

Modeling and Simulations of Quantum High Performance Computing

Travis S. Humble, Quantum Computing Institute, Oak Ridge National Laboratory

QAT Project: Heterogeneous Digital-Analog Quantum Dynamics Simulations

Pavel Lougovski, Oak Ridge National Laboratory

QCAT Project: Software Stack and Algorithms for Automating Quantum-Classical Computing

Pavel Lougovski, Oak Ridge National Laboratory; James Freericks, Georgetown University; Zohreh Davoudi, University of Maryland; Itay Hen, University of Southern California; Gregory Quiroz, Johns Hopkins University, Applied Physics Lab

Quantum Scientific Computing Open User Testbed (QSCOUT)

Peter Maunz, Andrew Landahl, Matthew G. Blain, Susan M. Clark, Daniel Lobser, Richard P. Muller
Melissa C. Revelle, Kenneth M. Rudinger, Christopher G. Yale, Sandia National Laboratories

Quantum Optimization and Learning and Simulation

Ojas Parekh, Sandia National Laboratories; Rolando Somma, Los Alamos National Laboratory; John Preskill, CalTech; Andrew Childs, University of Maryland

Methods and Interfaces for Quantum Acceleration of Scientific Applications (MIQASA)

Raphael Pooser, Oak Ridge National Laboratory; Ken Brown, Duke University; Sophia Economou, Virginia Tech

Quantum Computing Application Teams

Optimization, Verification, and Engineered Reliability of Quantum Computers (OVER-QC)

PI: Mohan Sarovar, Sandia National Laboratories

Team Members: Patrick Coles, Los Alamos National Laboratory; James Whitfield, Dartmouth College; Deepak Kapur, University of New Mexico

THE ADVANCED QUANTUM TESTBED

Irfan Siddiqi and Jonathan Carter, Lawrence Berkeley National Laboratory (LBNL)

William Oliver, Massachusetts Institute of Technology, Lincoln Laboratory (MIT-LL)

Quantum Performance Assessment

Robin Blume-Kohout, Erik Nielsen, Tim Proctor, Ken Rudinger, Kevin Young
Sandia National Laboratories (Albuquerque, NM)

Quantum processor technology is advancing at breakneck speed, with labs and companies announcing breakthrough devices seemingly every week. But it's hard to gauge how well any given device actually works. There aren't very many widely-accepted techniques for measuring performance, or even quantifying it in theory. Claims about hardware performance tend to be based on shallow metrics (e.g., number of qubits) or opaque ad-hoc benchmarks. Quantum hardware technology is outracing our ability to test devices and assess their performance.

Our goal, at Sandia's *Quantum Performance Laboratory*, is to provide solutions to this problem by leveraging our expertise and experience in quantum hardware characterization to develop a new and enabling toolbox of methods for testing and evaluating quantum hardware. The core of our program is *research*, to create new frameworks and techniques for performance assessment. But we also want to make our research *usable* by other scientists and engineers, both within and outside the DOE community. We will develop and deploy software tools that implement performance assessment protocols, which will allow quantum hardware developers and the US Government to assess quantum processors that they are developing or considering investing in. We also seek to partner with experimental researchers who are developing and improving qubit hardware, to test our methods and support their progress.

Our plan combines (1) informed R&D to create good metrics and benchmarks, (2) the development of reliable software to implement these techniques, and (3) engagement with experimental hardware developers and researchers to test and use the tools we develop. We intend to build on Sandia's demonstrated capability for characterization of 1- and 2-qubit gates using gate-set tomography (GST) and other protocols. We plan to extend this capability to comprehensive, holistic characterization of multiqubit quantum processors. We will develop protocols to measure the rates of specific errors, and we will determine how specific error modes impact a processor's overall performance. We will develop fair and well-motivated benchmarks to compare and contrast the performance of near-term "testbed-class" quantum processors. These benchmarks will incorporate kernels of promising DOE-relevant quantum algorithms for application-specific benchmarks. We will publish our methodologies, and release open-source implementations of them that can be used straightforwardly by experimentalists. We intend that these methods and their software implementations in the open-source *pyGSTi* package will support continued development of US quantum processors *and* help DOE to understand the performance of cutting-edge quantum hardware. We have already demonstrated this model at a small scale, by using *pyGSTi* to assess the error performance of 1-8 qubit devices. We plan to extend this capability to every form of "performance" that we can identify for testbed-class quantum processors.

We hope to partner with DOE-funded research teams in quantum algorithms to identify quantum algorithms that promise improved solutions to parts of DOE mission-relevant problems, map these algorithms to quantum circuits, and extract key subroutines or "kernels". We will then adapt these kernels to yield *test suites* of quantum circuits that can be run on near-term hardware, and capture its performance in real-world scenarios that are more relevant than how performance on random or contrived quantum circuits.

Efficient and Reliable Mapping of Quantum Computations Onto Realistic Architectures

Andrew Childs, University of Maryland (Principal Investigator)
Alexey Gorshkov, University of Maryland (Co-Investigator)
Michael Hicks, University of Maryland (Co-Investigator)
Xiaodi Wu, University of Maryland (Co-Investigator)
Robert Rand, University of Maryland (Key Personnel)

Quantum computers offer the prospect of efficiently solving certain problems that would be intractable for classical computers. Potential applications of quantum computers include several challenges of relevance to the DOE science and energy mission, including simulations of quantum mechanical systems, quantum approaches to approximate optimization, and numerical methods for scientific computing. However, the limited abilities of current quantum hardware make it challenging to realize these benefits with existing devices. The proposed work will develop methods and tools for implementing quantum algorithms as efficiently and reliably as possible on realistic hardware and for evaluating the capabilities of such devices, facilitating computation on near-term quantum processors.

Algorithms for quantum computers are typically formulated in an abstract circuit model that allows gates between any two qubits at the same cost. However, realistic quantum processors come with strong constraints on the possible interactions. Mapping a given quantum circuit onto such a realistic device will necessarily incur overhead depending on factors including the system architecture, the underlying quantum computation, and the method used to construct the mapping. Our proposed work aims to develop theoretical and computational tools for addressing this challenge that are efficient, robust, and incur minimal overhead. We will apply these tools to address both research areas suggested in the call for proposals: studying the interplay between architectural and algorithmic choices and quantitatively understanding the potential for near-term quantum devices to address scientific challenges.

This work will be organized around three interrelated research thrusts:

Circuit layout. To understand the effect of device architecture on the ability of a quantum processor to carry out quantum computations, it is crucial to develop methods for translating quantum circuits for algorithms of interest into ones that respect a given set of architectural constraints. A key thrust of our proposal is to produce a suite of algorithms for mapping quantum circuits to realistic architectures and implement them in software. Using these tools, we will explore the impact of architectural features on our ability to perform practical quantum computations, providing insights into the design of quantum architectures and evaluating the ability of realistic devices to carry out computations relevant to the DOE science mission.

Effect of interaction range. One of the defining features of any quantum computing platform is the range of available two-qubit interactions, with longer-range interactions allowing for faster computation. Some of the most advanced quantum computing platforms include long-range interactions. The second key thrust of our proposal is to harness such interactions to achieve computational speedups relative to architectures with nearest-neighbor interactions. We will also prove bounds on how quickly a given quantum algorithm can run using interactions of a given range. By combining the protocols and bounds that we develop, we will attempt to determine algorithm implementations for a given architecture that have the shortest possible running time, illuminating the relationship between device architecture and algorithm performance.

Formal methods. The final thrust of our work is the application of formal methods and techniques from the field of programming languages. Our goal is to develop novel tools to ensure that quantum programs are implemented and analyzed properly. Such tools will add confidence to the conclusions of our own analysis of near-term architectures and will also be of general use to the broader community. The result should be a much greater ability to reliably reason about the complexity and performance of near-term quantum devices.

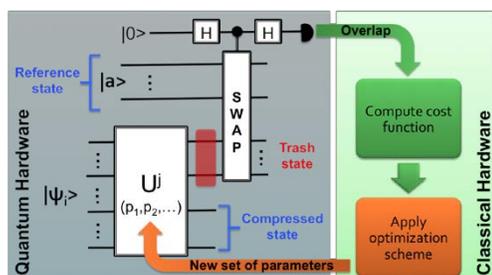
QAT4Chem

Quantum Algorithms, Mathematics and Compilation Tools for Chemical Sciences

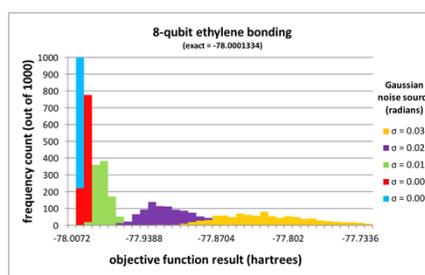
LBNL: Bert de Jong (Lead); ANL: Stefan Wild; University of Toronto: Alán Aspuru-Guzik

The Quantum Algorithms Team led out of Lawrence Berkeley National Laboratory is an integrated team of quantum algorithm developers, mathematicians, and computer scientists with a *mission to deliver algorithmic, computational and mathematical advances to enable scientific discovery in chemical sciences on quantum computers*. Our focus is on quantum chemistry simulations, which are an early exemplar of quantum computing, demonstrating the potential of various types of quantum devices to aid in scientific discovery in the chemical sciences.

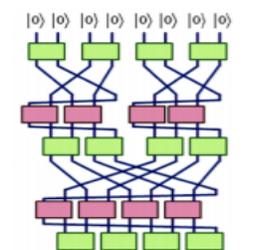
Our team has assessed the robustness of state preparation (<https://arxiv.org/abs/1809.05523>), devised new sparse techniques with lower gate depth (<https://arxiv.org/abs/1810.02327>), and have explored the use of tensor networks (<https://arxiv.org/abs/1803.11537>). Preparation of scientifically relevant states is a key challenge for chemical simulations. We have developed a protocol to discriminate between decoherence and information scrambling as a way to verify quantum circuits, which was experimentally validated on an ion trap (<https://arxiv.org/abs/1807.09087>). Our team is exploring a quantum autoencoder and its strongly reduced latent space for a quantum machine learning model. Gate depth of a quantum circuit plays a critical role in achieving results with high fidelity. The QAT team, in partnership with NASA, has demonstrated low depth circuits for k -local gates in QAOA and Trotterized fermionic Hamiltonians. We extended the open-source quantum simulation software framework ProjectQ with gate-level noise injection capabilities for error analysis, and used it to demonstrate that a rudimentary approach to mitigate errors in a CNOT gate operation can improve the gate fidelity. Additional efforts for error mitigation are ongoing. We have developed a large suite of stochastic classical optimizers, needed for variational quantum eigensolvers and qubit gate optimization. The optimizers have been packaged (github.com/scikit-quant) for integration in quantum computing software stacks. We are continuing to develop new and better stochastic optimizers.



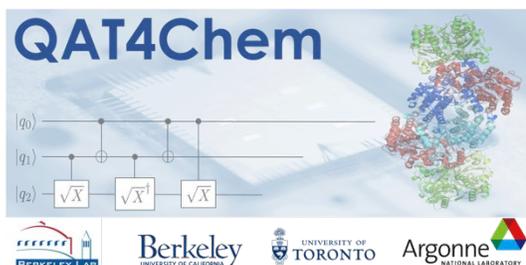
Quantum encoder compressing data



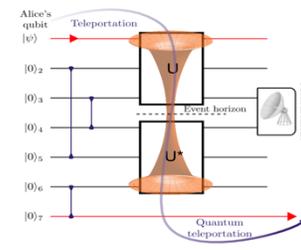
Role of noise on scientific results



Tensor networks to
encode wave functions



QAT4Chem.lbl.gov



Distinguishing between info
scrambling and decoherence

Superconducting Quantum Hardware Architecture Pathfinder

Jonathan L DuBois, Eric T. Holland, Yaniv J. Rosen, Luis A. Martinez, Xian Wu
Quantum Coherent Device Physics Group, Lawrence Livermore National Lab

The LLNL quantum testbed pathfinder effort is focused on exploration, design and demonstration of QC hardware architectures optimized for near term science applications. The computational unit cell of our current architecture, the “busmon”, consists of multiple long-lived cavity resonance modes coupled together with a transmon acting as a bus. Our approach is motivated by the need to address several challenges faced by current designs including hardware complexity, low connectivity of qubits, inefficient classical control, and complex fabrication requirements. Some of the advantages of the busmon are outlined below. Cavity states routinely demonstrate record levels of coherence with signal photon lifetimes of tens of milliseconds. Qudits encoded in cavity photon states are also efficient for storing and manipulating quantum information as the size of the effective computational Hilbert space is limited only by the total number of photons that can be controllably accessed and measured in each mode. The delocalized nature of quantum states in coupled cavity arrays naturally allows for a higher level of connectivity than conventional planar arrays of superconducting qubits resulting in faster, more complex entangling operations across the available state space. The ability to achieve accurate, minimum time, quantum gates for arbitrary unitary operations is a direct consequence of classical control amplitude, bandwidth and fidelity. In the LLNL approach a single, numerical waveform optimization combined with off-the-shelf, broadband arbitrary waveform generator enables n-body gate operations at or near a quantum speed limit allowing for fast and accurate entangling operations. Finally, we note that fabrication of the busmon is achieved with well-established and highly reproducible processes. Specifically, a single step aluminum liftoff process for the transmon devices and conventional high-precision machining for the cavities. The simplicity of this approach allows us to focus on device control and performance with minimal development of advanced fabrication techniques.

In the first year of the ASCR quantum testbed pathfinder program we demonstrated design, fabrication, characterization and universal control of a busmon with 6 effective qubits. The FY19 ASCR effort is focused on demonstrating the second generation of the busmon architecture with an effective computational Hilbert space equivalent to twenty all-to-all connected qubits.

This material is based upon work partially supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under the Quantum Testbed Pathfinder program and under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

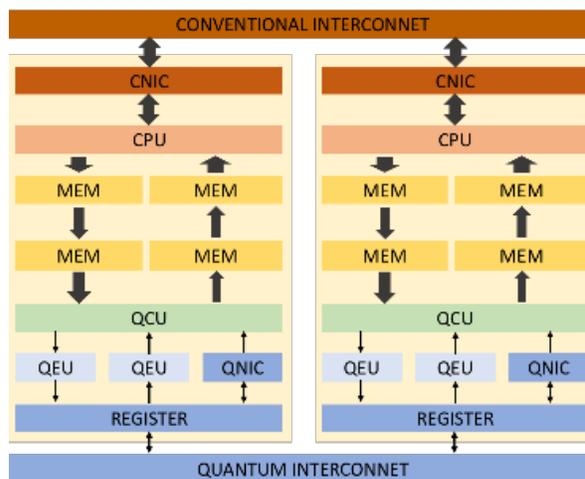
Modeling and Simulations of Quantum High Performance Computing

Travis S. Humble, *Quantum Computing Institute, Oak Ridge National Laboratory*

We use modeling and simulation to characterize the behavior of quantum high-performance computing (HPC) systems. These calculations provide estimates for the time, energy, complexity, and accuracy of both near-term and future quantum HPC. By modeling the integration of novel quantum processing units (QPUs) with conventional computer architectures, we are able to identify technological bottlenecks and recommend scientific advances necessary for realizing the greatest payoff from quantum computing. Our results indicate that existing massively, parallel processing architectures will require quantum networks to achieve optimal performance, while stand-alone, special purpose machines must be judiciously co-designed for specific applications.

Methodology

Modeling and simulation offer forward looking capabilities for assessing the impact of quantum computing on future scientific infrastructure. We develop models for quantum HPC as a hierarchy of computational devices interconnected by memory systems and networks. Our executable models include the data and instructions for quantum operations as well as the system logic for managing noisy and faulty computational resources. We simulate the operation of the conventional and quantum instructions using a combination of discrete-event and continuous-time simulation methods. We test examples that exercise the complete execution model for both near-term and future fault-tolerant QPUs.



Hierarchical model of quantum HPC system configured for quantum interconnected, massively parallel processing.

Key Insights

A notable insight from our results is that the performance of current NISQ QPUs operating in sampling mode is undercut by the bandwidth and latency of the calling process. Multiplexing and pipelining instructions lack natural analogs in the quantum domain and this manifests as a technological bottleneck for future architecture designs. Overcoming this bottleneck with traditional, multi-node parallelism is a temporary solution as the best quantum algorithms require coherent quantum memory across QPUs. Coincidentally, a quantum interconnect appears necessary to maintain quantum memory coherency, but this raises important questions about the speed with which classical control information can be distributed within the system.

Representative Publications

- High-performance Computing with Quantum Processing Units, ACM JETC (2017)
- Simulated Execution of Hybrid Quantum Computing Systems, Quantum Information Science, Sensing, and Computation X (2018)
- A Language and Hardware Independent Approach to Quantum–classical Computing, SoftwareX (2018)
- Consumer Applications of Quantum Computing, IEEE Consumer Electronics Magazine (2018)

This work supported by the Department of Energy, Office of Science, Early Career Research Program.

Heterogeneous Digital-Analog Quantum Dynamics Simulations

<https://hdaqds.ornl.gov/>

ORNL: Pavel Lougovski (Lead); University of Washington: Martin Savage

Enabling simulations of large-scale many-body systems is a long-standing problem in scientific computing. Recent progress in designing digital and analog quantum simulation algorithms hints at the exciting possibility of performing simulations that are beyond the reach of all existing and future classical supercomputers. Despite the tremendous progress, there is still a gap between the resources required by state-of-the-art quantum algorithms and the resources offered by available and near-future quantum hardware.

The main goal of this project is to bridge the gap between quantum algorithms and hardware by developing a novel algorithmic approach to quantum simulations. Our approach combines individual algorithmic strengths of analog and digital quantum simulations and it is heterogeneous across quantum hardware, classical and quantum algorithms. For specific scientific domain applications, we focus on designing hardware-efficient quantum algorithms for simulating the dynamics of nuclear, strongly-correlated electron, and lattice gauge quantum field systems.

To accomplish our objective, we have assembled a multidisciplinary team of computer scientists, applied mathematicians, scientific application domain experts, and quantum computing researchers. As part of this project our team develops and analyzes quantum simulation algorithms for computing: i) the binding energy of light nuclei; ii) spin lattice systems dynamics; iii) lattice gauge quantum field dynamics.

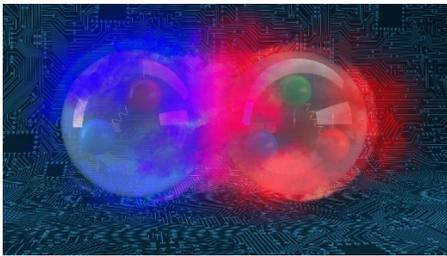


Figure 1. Quantum simulations of the deuteron
[Phys. Rev. Lett. **120** 210501 (2018)]

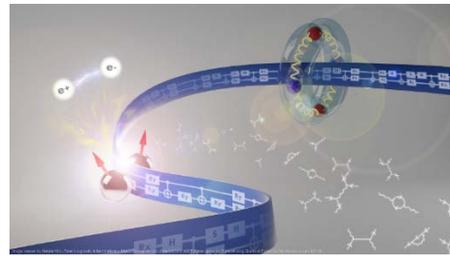


Figure 2. Quantum simulations of the Schwinger model
[Phys. Rev. A **98** 032331 (2018)]

Project Publications

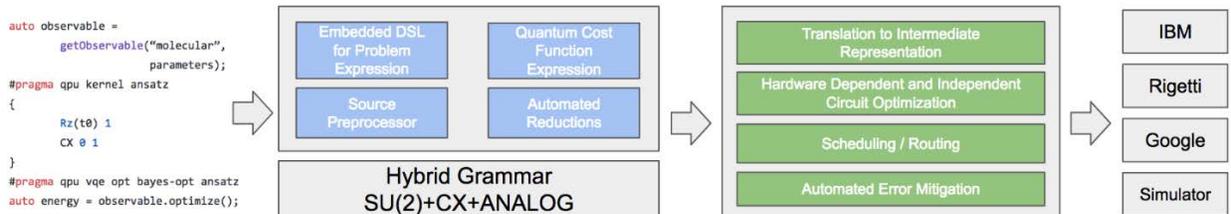
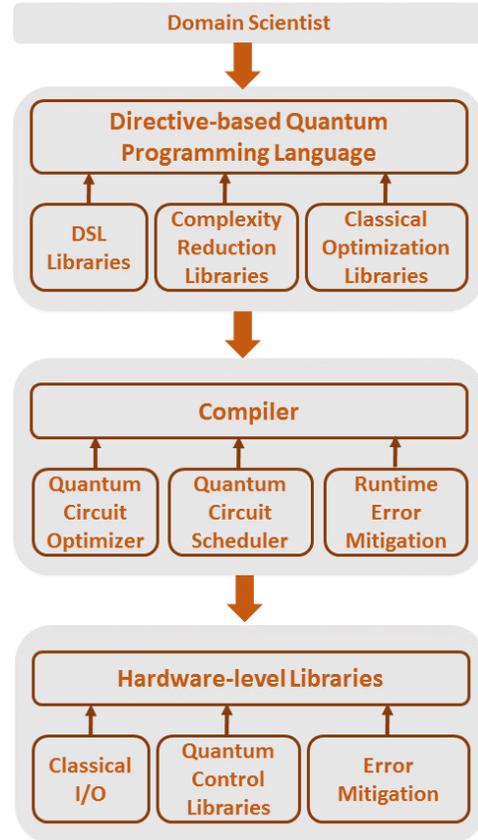
- N. Klco *et al.*, “Quantum-classical computation of Schwinger model dynamics using quantum computers”, Phys. Rev. A. **98**, 032331 (2018)
- E. F. Dumitrescu *et al.*, “Cloud Quantum Computing of an Atomic Nucleus”, Phys. Rev. Lett. **120**, 210501 (2018)
- L. Lamata *et al.*, “Digital-Analog Quantum Simulations with Superconducting Circuits”, Advances in Physics: X **3**, 1457981 (2018)
- A. J. McCaskey *et al.*, “A language and hardware independent approach to quantum-classical computing”, Software X **7**, 245-254 (2018)
- Y. Ding *et al.*, “Experimental Implementation of a Quantum Autoencoder via Quantum Adders” arXiv:1807.10643 (2018)
- N. Klco, M. Savage, “Digitization of Scalar Fields for NISQ-Era Quantum Computing.” arXiv:1808.10378 (2018)
- H.-H. Lu *et al.*, “Simulations of Subatomic Many-Body Physics on a Quantum Frequency Processor” arXiv:1810.03959 (2018)

Software Stack and Algorithms for Automating Quantum-Classical Computing

ORNL: Pavel Lougovski (Lead); Georgetown University: James Freericks; UMD: Zohreh Davoudi; USC: Itay Hen; JHU Applied Physics Lab: Gregory Quiroz

For a long time, developing real-world applications for quantum computing (QC) has been a prerogative of a select few quantum information specialists with access to experimental hardware. Cloud QC pioneered by IBM and Rigetti, as well as the ongoing efforts by DOE ASCR to create testbeds, are poised to change this predisposition and bring QC within the reach of domain scientists. However, to create and execute a quantum simulation algorithm, software tools are needed that address multiple aspects of QC : (i) translation of the problem at hand from a domain-specific language (DSL) into the language of qubits and quantum circuits, (ii) optimization of the resulting circuits to meet hardware constraints, (iii) implementation of rudimentary error mitigation strategies at runtime, and (iv) postprocessing of measurement results. As a result, the wide adoption of QC by the research community is complicated by *the problem of translating scientific applications into functional quantum programs that return accurate results when executed on real-world faulty quantum hardware.*

This project addresses this problem by developing an open-source algorithm and software stack that will automate the process of designing, executing, and analyzing the results of quantum algorithms, thus enabling new discovery across many scientific domains. Prompted by limitations on algorithms and models of computation imposed by available and near-term quantum hardware – noisy quantum gates, limited qubit connectivity, and short circuit depth to name a few – we focus the stack on implementing a hybrid quantum-classical model of computation for various types of quantum hardware with particular emphasis on scientific applications in quantum field theory, nuclear physics, condensed matter, and quantum machine learning.



```

auto observable =
  getObservable("molecular",
               parameters);
#pragma gpu kernel ansatz
{
  Rz(theta) 1
  CX 0 1
}
#pragma gpu vqe opt bayes-opt ansatz
auto energy = observable.optimize();
    
```

Quantum Scientific Computing Open User Testbed (QSCOUT)

Peter Maunz, Andrew Landahl, Matthew G. Blain, Susan M. Clark, Daniel Lobser, Richard P. Muller
Melissa C. Revelle, Kenneth M. Rudinger, Christopher G. Yale,

Sandia National Laboratories, Albuquerque, NM

This project will develop the Quantum Scientific Computing Open User Testbed (QSCOUT), a testbed laboratory to assess the potential of near-term quantum hardware to address scientific computing applications of interest to the U.S. Department of Energy (DOE) and its Advanced Scientific Computing Research (ASCR) program.

We will realize QSCOUT as a trapped-ion system to leverage its unique architectural advantages. Trapped-ion systems enable highly connected qubit layouts (up to all-to-all connections), they feature a reconfigurable qubit arrangement that can be adapted to different algorithms, they offer the highest fidelity single- and two-qubit gates, they admit multiple ways to scale to large processors, and there is a path for them to execute coherent operations networked beyond a single processor. All-to-all connectivity can significantly reduce the number of gates required to execute a quantum circuit. The reconfigurability of the system has important consequences for algorithm and software design. These advantages also enable trapped-ion processors to emulate other quantum processors, allowing one to explore the impacts limited qubit connectivity has on other platforms. High fidelity operations enable larger algorithms without the need to employ error correcting codes.

The low-level access needed for external researchers to understand how to use these devices and the ability to modify these devices and experiment with different implementations of gates and circuits requires an open-hardware approach. Therefore, we will realize the testbed system at Sandia by leveraging, extending, and refining existing quantum processor designs. Furthermore, we believe that operating the processor at cryogenic temperatures is essential to realize sufficient operation fidelities and system uptimes. We will make a room-temperature and a cryogenic system available to external researchers at the end of year 1 and 4, respectively. Sandia is a world-leader in the design and fabrication of scalable surface-electrode ion traps and has previously demonstrated high-fidelity single- and two-qubit gates in microfabricated surface traps. We will couple our high-fidelity hardware with correspondingly powerful software capabilities that allow full control over the signals, gates, and circuits. This will allow users to fully explore how well quantum hardware can address scientific applications of interest to DOE/ASCR.

Being a small-scale system, QSCOUT's purpose is not to demonstrate "quantum advantage," *i.e.*, solve a mission-relevant problem faster or better than a classical supercomputer. Rather, QSCOUT's purpose is to enable external researchers to explore and evaluate quantum algorithms for some of those mission-relevant problems. QSCOUT will provide unrestricted access to the internals of the quantum hardware and enable users to adapt and modify the quantum gates and underlying pulse sequences. We will support and advise users in the functionality of the QSCOUT testbeds and provide users with exemplar programs demonstrating the most relevant scientific applications, including the variational quantum eigensolver that determines the electronic structure of materials and molecules.

QSCOUT and its access to all relevant implementation details will allow DOE users to *understand the potential* of quantum platforms for scientific inquiries and it will allow them to take the first steps toward *realizing the potential* of high-fidelity quantum computing.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.



[qoalas.sandia.gov]

QOALAS

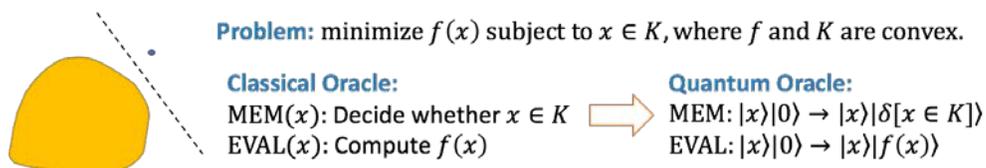
Quantum Optimization and Learning and Simulation

SNL: Ojas Parekh (Lead); LANL: Rolando Somma; CalTech: John Preskill;

UMD: Andrew Childs

The QOALAS Quantum Algorithms Team explores the impact of quantum computing in three interrelated areas underpinning the DOE mission: quantum simulation, optimization, and machine learning. Our work leverages connections among these areas to fuel new applications of quantum information processing to science and technology. We are an interdisciplinary team of theoretical physicists and theoretical computer scientists from Caltech, Los Alamos Lab, Microsoft Research, NIST, Sandia National Labs, and the University of Maryland.

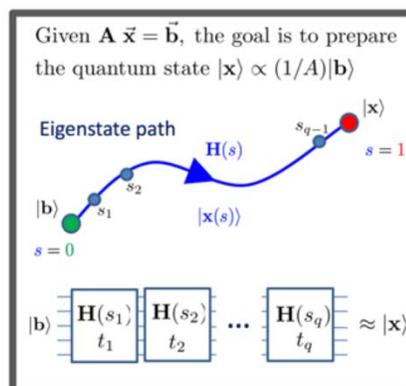
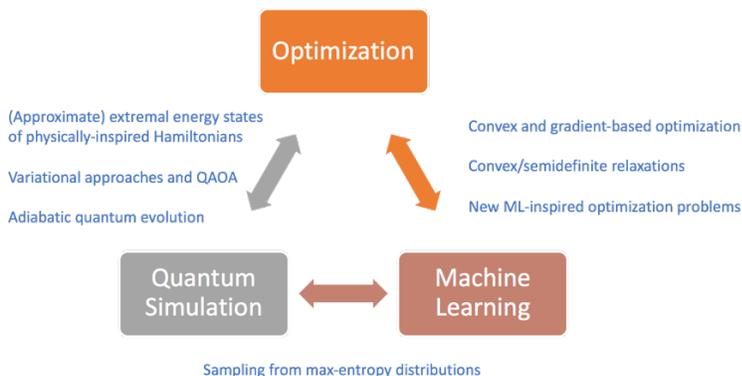
Quantum Simulation. Quantum computers offer a natural approach to efficiently simulating the dynamics of quantum systems, a fundamental physical problem intractable for classical computers. We have produced stronger bounds on the performance of quantum simulations based on product formulas, through a novel use of randomization [arXiv:1805.08385]. We are also investigating applying quantum computation to problems in high-energy physics, where much work remains to develop simulations of general field theories.



New quantum algorithm for convex optimization using $\tilde{O}(n)$ oracle queries vs $\tilde{O}(n^2)$ classically [arXiv:1809.01731]

Optimization. We seek to understand how quantum resources may be used to reap gains in speed or solution quality over current classical approaches. We have developed new quantum techniques for rigorously approximating ground states of physically motivated quantum Hamiltonians. Our work connects widely used techniques in discrete optimization to central topics in condensed matter physics. We have also recently developed some of the first quantum algorithms for convex optimization (see figure above).

Machine learning. Linear algebraic kernels are a fundamental ingredient in machine learning and serve as a building block for advanced quantum machine learning algorithms. We have designed a new quantum algorithm for solving linear systems, inspired by adiabatic quantum computing, that is much simpler than previous approaches with comparable performance (see figure below).



A simpler quantum linear system solver [arXiv:1805.10549]

Methods and Interfaces for Quantum Acceleration of Scientific Applications (MIQASA)

Raphael Pooser (ORNL), Ken Brown (Duke University), Sophia Economou (Virginia Tech)

The MIQASA project explores the design and development of quantum computing for scientific applications in chemistry, nuclear physics, quantum field theory, and machine learning. The project has several major goals: 1) To determine the feasibility of running scientific applications of interest to ASCR and DOE on near-term quantum computer (QC) hardware. 2) To foster industry relationships and harness multiple QCs for DOE applications, and to benchmark them. This goal requires the development of both benchmarks and suitable metrics that can accurately communicate to the community the utility or potential utility of these machines. Providing accurate benchmarks and metrics also helps hardware developers understand how to target future performance for scientific applications, thus seeding the codesign loop. Our project effectively delivers a prototype example for future codesign methods for QCs. 3) To foster academic partnerships and engage in educational outreach. We achieve this goal by partnering with universities, speaking in public forums, and maintaining a set of resources for the community such as specification sheets for quantum processors that help researchers understand the metrics we derive under this project.

We achieve our first two goals by benchmarking quantum computers that are available today and by gathering and organizing metrics into an easy to understand format. Benchmarks consist of traditional tomography to assess quantum gate performance, larger collections of gates that represent computational primitives, and the applications themselves, which are suitable benchmarks in this early, pre-fault tolerant era of quantum computing, often referred to as noisy-intermediate-scale quantum (NISQ). In order to run these benchmarks and applications, we also study how to efficiently program quantum computers. We leverage ORNL's hardware-agnostic programming environment to automate tasks such as error mitigation, qubit tapering, variational optimization, and qubit selection and placement.

Examples of benchmarks from each of the above categories are:

Gate constructions and tomography, including the design of waveforms to produce a CPHASE gate: We simulate RF pulse controls for transmon qubits using Virginia Tech's simulator and design entangling gates using the idea of Speeding up Waveforms by Inducing Phases to Harmful Transitions (SWIPHT) in order to produce faster two-qubit gates, and subsequently test these on real hardware. We benchmark existing gate constructions against new gates with tomography and show that new CZ, and CNOT achieve more than 99% fidelity. Our project is integrating these gate constructions into higher level algorithms on hardware using OpenPulse in partnership with IBM.

Algorithmic Primitives: We use machine learning to train quantum registers to prepare specific wavefunctions and encode information in probability amplitudes. Our machine learning benchmarks test a machine's capability to perform basic primitive operations hybrid algorithms. We use these benchmarks to determine optimal circuits that take into account the hardware's strengths. We also produce benchmarks which test the computational power of random gate constructions, which provides a supremacy benchmark.

Applications: We test multiple scientific applications in order to determine the suitability for QCs to run them now or in the future, and to determine what, if any, quantum advantage is to be had. In some cases the advantage may come from accuracy rather than execution speed, for example. Examples include estimating the bound state of the deuteron and other nuclei; estimating time evolution or symmetry phases of the Schwinger model or interacting scalar field theories; and an array of alkali-hydride molecular ground state calculations for a chemistry benchmarking suite.

When taken in aggregate, these benchmark categories paint a holistic picture of a QC's performance potential. When benchmarks represent scalable algorithm implementations, and gate benchmarks indicate low noise, high performance is an indicator of scalable performance in the fault-tolerant, beyond-NISQ era as well. When performance is low, our benchmarks indicate key performance metrics for next generation devices to improve in order to reach scalable implementations. The end result is a determination of the feasibility to run scientific applications on current and near-future hardware in the NISQ era and beyond.

Optimization, Verification, and Engineered Reliability of Quantum Computers (OVER-QC)

SNL: Mohan Sarovar (Lead); LANL: Patrick Coles; Dartmouth University: James Whitfield; UNM: Deepak Kapur

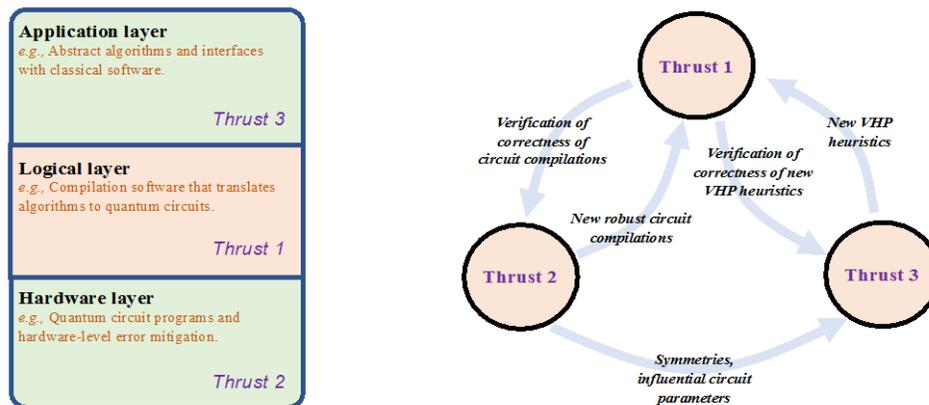
Quantum technologies, especially quantum computers, show great promise for revolutionizing high-performance computing and simulation. As prototype quantum computers come online, it is becoming clear that obtaining useful output from such devices will require layers of sophisticated classical software that provide interpretation and analysis of the quantum computer's state and output. We will develop critical components in this “software stack” with a particular focus on enabling near-term Noisy Intermediate-Scale Quantum (NISQ) technologies. We have identified three critical needs for near-term quantum computing platforms, and the project is structured around three thrusts that address these needs:

Thrust 1: develop capabilities to verify and certify translations of abstract quantum algorithms to quantum circuit programs;

Thrust 2: develop tools to identify the reliable information that can be extracted from near-term devices in the absence of fault-tolerant operation;

Thrust 3: develop interfaces for variational hybrid quantum-classical processors (VHPs) that enable applications by connecting classical algorithms to quantum co-processors.

The following figure illustrates how these thrusts fit into the conventional picture of a quantum computing software stack, and how the outcomes of each thrust inform the other ones.



A key aspect of the proposed research is that we will not only develop software tools targeted for use with near-term quantum computers, but we will also invest significant effort into understanding the reliability and computational power of near-term devices. In particular, we will address two key questions, (1) is it possible to engineer the dynamics of NISQ devices so that they have robust algorithmic properties in the presence relevant error models, and (2) what is the real computational advantage posed by algorithms running on NISQ devices, and specifically, what level/type of noise permits some computational advantage over classical algorithms and heuristics for the same tasks. Answering these questions will not only guide the development of quantum algorithms tailored to near-term quantum devices, but it will also reveal the fundamental computational potential of non-fault-tolerant quantum devices.

THE ADVANCED QUANTUM TESTBED

Irfan Siddiqi and Jonathan Carter
Lawrence Berkeley National Laboratory (LBNL)

William Oliver
Massachusetts Institute of Technology, Lincoln Laboratory (MIT-LL)

The Advanced Quantum Testbed (AQT) is a collaborative research facility for the advancement of quantum computation based on superconducting circuits. We will develop and implement quantum algorithms co-designed for noisy, intermediate scale quantum hardware with a focus on computational problems relevant to the scientific computing mission of the DOE, including applications in optimization, materials science, and cosmology. Our technology strategy maintains the state-of-the-art by drawing from multiple hardware and software resources, and by transitioning specific advances in basic quantum information science research programs (and fledgling industry) to the testbed where quantum algorithms can then be executed and refined. We aim to move quantum algorithms beyond single, academic, proof-of concept demonstrations to a broad range of informative and impactful calculations. Testbed users will have full access to the hardware (including detailed data concerning architecture, operation, limitations, and imperfections), participate in its evolution, and share results to maximize the utility of nascent quantum hardware. An on-site catalyst team, expert in quantum hardware and software, will advance the basic science mission of the testbed and interact with a broad community of users to help them make effective use of our resources.

The AQT will be a unique facility for quantum computation, bridging a key gap between highly exploratory academic research efforts and highly specialized instances of commercially available, cloud-based and physical resources. The overall facility vision revolves around a flexible architecture for deploying state-of-the-art superconducting quantum technology, including processors as well as control systems & software. Distinct from academic and industrial players, the AQT is a priori approach agnostic, not bandwidth limited, and has a broad mandate to upgrade all aspects of facility operation through a combination of user influenced internal development and strategic partnerships with external entities. The latter will permit technology rooted outside of the national laboratory environment, for example mature designs in the academic world as well evolving industrial products, to be deployed and field-tested, thereby providing valuable feedback to these other inhabitants of the quantum eco-system and also providing facility users with access to the latest pre-competitive quantum devices. The facility will therefore play a central role in benchmarking different types of quantum technology using a common hardware platform, naturally supporting the future development of standardized metrics.