

# Assessing and Advancing the Potential of Quantum Computing: A NASA Case Study

Eleanor G. Rieffel<sup>a,\*</sup>, Ata Akbari Asanjan<sup>d,c</sup>, M. Sohaib Alam<sup>a,c</sup>, Namit Anand<sup>a,b</sup>, David E. Bernal Neira<sup>a,c</sup>, Sophie Block<sup>a,c</sup>, Lucas T. Brady<sup>a,b</sup>, Steve Cotton<sup>a,b</sup>, Zoe Gonzalez Izquierdo<sup>a,c</sup>, Shon Grabbe<sup>a</sup>, Erik Gustafson<sup>a,c</sup>, Stuart Hadfield<sup>a,c</sup>, P. Aaron Lott<sup>a,c</sup>, Filip B. Maciejewski<sup>a,c</sup>, Salvatore Mandrà<sup>a,b</sup>, Jeffrey Marshall<sup>a,c</sup>, Gianni Mossi<sup>a,b</sup>, Humberto Munoz Bauza<sup>a,b</sup>, Jason Saied<sup>a</sup>, Nishchay Suri<sup>a,c</sup>, Davide Venturelli<sup>a,c</sup>, Zhihui Wang<sup>a,c</sup>, Rupak Biswas<sup>a</sup>

<sup>a</sup>Quantum Artificial Intelligence Lab. (QuAIL), Exploration Technology Directorate, NASA Ames Research Center, Moffett Field, CA, 94035, USA

<sup>b</sup>KBR, 601 Jefferson St., Houston, TX, 77002, USA

<sup>c</sup>USRA Research Institute for Advanced Computer Science (RIACS), Mountain View, CA, 94043, USA

<sup>d</sup>Data Science Group (DSG), Exploration Technology Directorate, NASA Ames Research Center, Moffett Field, CA, 94035, USA

## Abstract

Quantum computing is one of the most enticing computational paradigms with the potential to revolutionize diverse areas of future-generation computational systems. While quantum computing hardware has advanced rapidly, from tiny laboratory experiments to quantum chips that can outperform even the largest supercomputers on specialized computational tasks, these noisy-intermediate scale quantum (NISQ) processors are still too small and non-robust to be directly useful for any real-world applications. In this paper, we describe NASA's work in assessing and advancing the potential of quantum computing. We discuss advances in algorithms, both near- and longer-term, and the results of our explorations on current hardware as well as with simulations, including illustrating the benefits of algorithm-hardware co-design in the NISQ era. This work also includes physics-inspired classical algorithms that can be used at application scale today. We discuss innovative tools supporting the assessment and advancement of quantum computing and describe improved methods for simulating quantum systems of various types on high-performance computing systems that incorporate realistic error models. We provide an overview of recent methods for benchmarking, evaluating, and characterizing quantum hardware for error mitigation, as well as insights into fundamental quantum physics that can be harnessed for computational purposes.

**Keywords:** quantum computing, NASA

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>	<b>3.3</b>	<b>Advancements in Quantum-Inspired Generative Models: RBMs in VAEs and Invertible Flows</b>	<b>11</b>
<b>2</b>	<b>Quantum Optimization Algorithms and Sampling</b>	<b>3</b>	<b>3.4</b>	<b>Quantum-Compatible Discrete Deep Generative Models</b>	<b>11</b>
2.1	QAOA Analysis and Development	3	<b>3.5</b>	<b>Quantum-Enhanced Optimization of Quantum Heuristics</b>	<b>12</b>
2.2	Entanglement verification in QAOA	5	<b>4</b>	<b>Quantum Computing for Simulation</b>	<b>12</b>
2.3	Iterative Quantum Algorithms	5	4.1	Quantum Simulation for Condensed Matter Physics	12
2.4	Quantum Annealing Advances	5	4.1.1	Kitaev Spin Models	12
2.5	Quantum Optimization Applications	7	4.1.2	Fermi-Hubbard Model	12
2.6	Distributed Quantum Computing	7	4.2	Quantum Computing for Material Discovery	13
2.7	Quantum Algorithm Design and Problem Encoding	8	4.2.1	Advancements in State Preparation Techniques	13
2.8	Physics Inspired Heuristics	8	4.2.2	Many-body Scar Preparation	13
<b>3</b>	<b>Quantum-compatible Machine Learning Algorithms</b>	<b>10</b>	4.2.3	Advancements in Hamiltonian Simulation Techniques	13
3.1	Quantum-Assisted Variational Autoencoder	10	4.3	Quantum Simulation for Chemistry	14
3.2	Optimizing Adversarial Networks Through Quantum Annealing	11	4.3.1	Advances in classical solvers for chemistry Hamiltonians	14
			4.4	Open Quantum System Simulation	14
			4.5	Quantum simulation of High Energy Physics	14

\*Corresponding author

Email address: eleanor.rieffel@nasa.gov (Eleanor G. Rieffel)

<b>5</b>	<b>Tools for Quantum Computing Investigations</b>	<b>15</b>
5.1	Classical Methods for Simulating Quantum Circuits . . . . .	15
5.2	Photonics . . . . .	16
5.3	Quantum Computing Benchmarking . . . . .	16
5.4	Noise Modeling and Characterization . . . . .	16
5.5	Quantum Error Correction and Mitigation . . . . .	17
5.6	Pulse Control . . . . .	17
5.7	Generating Hard Problem Instances . . . . .	18
5.8	Randomized Benchmarking . . . . .	18
5.9	Master Equation Emulation . . . . .	19
<b>6</b>	<b>Mechanisms for Quantum Computation</b>	<b>19</b>
6.1	Wigner Friend Scenario . . . . .	19
6.2	Information Scrambling . . . . .	19
6.3	Stoquasticity . . . . .	20
6.4	Relaxation Mechanisms in Quantum and Classical Transverse-Field Annealing . . . . .	20
6.5	Acceleration of Quantum Tunneling in Spin Systems . . . . .	20
6.6	Engineering Topology . . . . .	20
6.7	Compilation approaches . . . . .	20
6.8	Quantum Computing Architectures . . . . .	21
<b>7</b>	<b>Future Outlook</b>	<b>21</b>
<b>8</b>	<b>Acknowledgements</b>	<b>21</b>
<b>9</b>	<b>About the NASA QuAIL Team</b>	<b>22</b>

## 1. Introduction

We describe work by NASA’s Quantum Artificial Intelligence Laboratory (QuAIL) and collaborators on assessing algorithms and architectures across application areas targeted by NASA. This broad range of research activities includes developing fundamental theory, compilation and error mitigation, quantum machine learning, optimization, and simulation algorithms, as well as benchmarking and analysis of algorithms and hardware. Much of this work is motivated and enabled via close collaborations with industry, government, and academic groups building quantum hardware, and illustrates the benefits of algorithm-hardware co-design in the near-term Noisy-Intermediate Scale Quantum (NISQ) era [1]. This work builds on prior QuAIL team work [[2, 3]] and focuses on computational challenges in optimization, machine learning, simulations, tools supporting evaluation and algorithmic development, and theory and numerical work that deepens understanding of physical mechanisms that can be harnessed for quantum computation. For clarity, citations in double brackets, [[ ]], highlight works involving QuAIL team members, while citations in single brackets, [], refer to third-party works.

Throughout our work, hardware-algorithm co-design is always present, with algorithmic needs feeding into the hardware design, hardware capabilities, and limitations feeding into algorithm design, and hardware noise characterization feeding into error mitigation techniques and the development

of noise-aware algorithms. Key to our research has been open quantum system modeling, analysis, and simulation, leading to an improved understanding of physical implementation and co-design across multiple architectures, including gate-model, measurement-based, and quantum annealing (QA) architectures. Much of this work has been focused on available and proposed NISQ architectures, enabling empirical investigations for testing and developing theories and guiding new algorithm and hardware co-design approaches based on effective and efficient application-driven requirements to guide between where hardware stands and purely theoretically motivated architectures and algorithms. This work enabled a deeper understanding of the mechanisms underpinning quantum annealing, such as tunneling, thermalization, and many-body delocalization, and their interaction with embedding parameters and annealing schedules. Early work on annealing architectures has been extended to and built upon to analyze multiple gate-based architectures and algorithms.

As part of the NASA community, the QuAIL team has had access to domain experts with a variety of challenging computational problems. These interactions, and those with partners, have led to work developing benchmark instance sets capturing aspects of computational challenges arising in space, aeronautics, and earth science missions, as well as in the exploration of quantum, quantum-classical hybrid, and physics-inspired approaches to tackling these problems. Some example application areas include planning and scheduling, anomaly detection, fault diagnosis, autonomous air traffic management, machine learning, and simulation of materials, chemistry, and high energy physics. The QuAIL team has led and supported multiple benchmarking and analysis activities focused on evaluating quantum architectures and approaches, comparing quantum and quantum-classical hybrid algorithms against state-of-the-art classical approaches, and providing resource estimation for running such algorithms at application scale on future quantum architectures.

There are recurring themes and lessons learned from our recent work, which include:

- benefits of hardware-algorithm co-design in both the near-term NISQ era and the long-term fault tolerant era;
- usefulness of experimentation on NISQ devices, from error characterization to insights into quantum and quantum-classical hybrid algorithms in the fault-tolerant and application-scale regimes;
- physics-inspired classical algorithms can be used at application scale today;
- synergistic advances of classical and quantum machine learning (QML) research, including quantum-ready classical algorithms that can be used today but are ready for replacement of some classical subroutines with quantum subroutines once quantum hardware matures sufficiently;
- the use of classical high-performance computing (HPC) for validation and benchmarking of quantum devices, quantum simulations, and algorithm exploration;

- importance of and methods for generating hard problem instances for benchmarking quantum, quantum-inspired, and classical devices.

This paper is organized as follows. Sec. 2 surveys recent work on quantum algorithms for optimization and sampling. Quantum-compatible machine learning work is covered in Sec. 3. We discuss applications of quantum computing to simulating quantum systems, including quantum chemistry, materials, condensed matter, high energy, and fundamental physics in Sec. 4. Tools for supporting quantum computing investigations are described in Sec. 5. Lastly, Sec. 6 focuses on work associated with mechanisms for quantum computation.

## 2. Quantum Optimization Algorithms and Sampling

Quantum algorithm research can be broken into two broad categories: algorithms focused on near-term devices and applications and algorithms focused on long-term fault-tolerant computers. Near-term algorithm research enables demonstrations of quantum computing capabilities on current devices and the testing of error characterization, mitigation, and correction protocols within certain algorithms. Of particular interest is algorithmic work on near-term devices that can give insight into quantum algorithms that will be relevant in the fault-tolerant era. Hardware-algorithm co-design is of particular importance in the pre-fault-tolerant era and in the early fault-tolerant era when resources will be heavily constrained. To date, a number of quantum algorithms have been proposed for combinatorial optimization problems (e.g., [4, 5, 6, 7, 8, 9]). Not only does the size of current quantum hardware pose a challenge, but they are highly susceptible to various sources of error. Designing algorithms that can be run on the limited hardware of today and give insight into the design of quantum algorithms that will be effective in the fault-tolerant era is a challenge. Not only do the limitations and constraints of current quantum hardware influence algorithmic work, but algorithmic considerations also impact hardware designs and priorities. This cycle of co-design will first be illustrated in the setting of quantum optimization.

Quantum algorithms come in many different forms, from gate-based models to quantum annealing, and are diverse enough in scope to also include sampling methods and binary decision problems. Quantum computers are always hybrid in nature; they rely on classical controllers. Many of the most effective algorithms, even in the future, are expected to be hybrid, integrating communication between the quantum and classical processors during algorithm execution. Optimization will provide our first and most extensive look at hybrid and well as fully quantum algorithms. A broad recent overview of quantum optimization [[10]] examines potential opportunities and challenges across many approaches and application areas.

A new paper, [[11]], surveys HPC and quantum computing methods and perspectives for optimization.

### 2.1. QAOA Analysis and Development

The Quantum Approximate Optimization Algorithm (QAOA), a prominent near-term quantum optimization algorithm [5], was extended to the Quantum Alternating Operator

Ansatz by the QuAIL group [[12, 2]]. This algorithm uses a bang-bang quantum circuit that alternates between pulses related to the problem Hamiltonian and a mixer Hamiltonian. The pulse lengths, variational parameters optimized by an outer classical loop, aim to produce a state minimizing the problem Hamiltonian’s energy. Analytical parameter determination remains an open question, with ongoing advancements including hardware-algorithm co-design strategies that can extract the full potential of currently available quantum devices.

In the cited work [[13]], advancements in parameter setting strategies for QAOA were introduced, focusing on constraint satisfaction problems with polynomial growth in distinct cost values. The model, rooted in empirical QAOA observations, demonstrated effectiveness in the numerical evaluation of the heuristic for MaxCut on unweighted Erdős-Rényi random graphs, finding that for three levels of QAOA, the heuristic matches approximation ratios previously achieved with globally optimized methods. Additionally, for levels up to 20, the heuristic showed parameters with monotonically increasing approximation ratios, indicating its scalability and potential efficacy in deeper QAOA applications.

Despite much effort, obtaining rigorous performance bounds for quantum algorithms such as QAOA beyond relatively basic cases remains challenging. In [[14]], analytical results were obtained for MaxCut problems on graphs utilizing the lowest depth realization of QAOA. It was shown that the expected value of each term in the cost function, corresponding to the graph edges, depends only on the local neighborhood of each edge. This property extends to deeper QAOA circuits but over larger neighborhoods. Similar notions of local algorithms exist classically.

In the recent work of [[15]], the power of both quantum and classical local algorithms for approximately solving Max kXOR, a natural generalization of MaxCut from graphs to hypergraphs was considered. In Max kXOR, each constraint is the XOR of exactly  $k$  variables and a parity bit. On instances with either random signs (parities) or no overlapping clauses and  $D+1$  clauses per variable. This is analogous to triangle-free and regular graph instances of MaxCut, respectively. The expected satisfying fraction of the depth-1 QAOA is exactly calculated and compared with a generalization of the classical local threshold algorithm from [16]. Notably, the quantum algorithm outperforms the threshold algorithm for  $k > 4$ . On the other hand, potential difficulties for QAOA to achieve computational quantum advantage on this problem are highlighted. We first compute a tight upper bound on the maximum satisfying fraction of nearly all large random regular Max kXOR instances by numerically calculating the ground state energy density  $P(k)$  of a mean-field  $k$ -spin glass [17]. The upper bound grows with  $k$  much faster than the performance of both local algorithms. A new obstruction result for low-depth quantum circuits (including QAOA), when  $k = 3$ , generalizing a result of [18] for the  $k = 2$  case, is also identified, and it is conjectured that a similar obstruction exists for all  $k$ . Such obstruction results are important steps towards a better understanding of the power of quantum circuits for optimization problems.

Several different QAOA approaches to solving the combi-

natorial circuit fault diagnostic (CCFD) problem are introduced and analyzed in [19]. Comparing these approaches on a family of dense and highly connected circuits supports the intuition that approaches that are more closely tailored to exploiting the structure of the underlying optimization problems can have better performance than more general approaches.

In [20], a framework for analyzing layered quantum algorithms such as quantum alternating operator ansätze is developed. The framework relates quantum cost gradient operators, derived from the cost and mixing Hamiltonians, to classical cost difference functions that reflect the cost function neighborhood structure of the problem at hand. Exact general expressions for expectation values are derived as series expansions in the algorithm parameters, cost gradient operators, and cost difference functions. This enables novel interpretability and insight into QAOA behavior in various parameter regimes, including but not limited to small-parameter regimes. For single-level QAOA1, it is shown that the leading-order changes in the output probabilities and cost expectation value explicitly for arbitrary cost functions, demonstrating that, for sufficiently small positive parameters, probability flows from lower to higher cost states on average. This result is leveraged as it is shown that QAOA always beats random guessing. Several example applications of the framework, including Quadratic Unconstrained Binary Optimization (QUBO) problems and variants of MaxSat, are also considered, as well as the extension to mixing unitaries beyond the transverse-field mixer for constrained optimization problems [12].

To better understand the behavior of QAOA, an approach to analyzing deep circuits with gradually varying unitaries is introduced in [21] and note several surprising phenomena. First, the ground state of the mixer Hamiltonian directly connecting to a highly excited state of the cost Hamiltonian will result in poor QAOA performance even in the adiabatic limit. Secondly, shallower circuits are shown to outperform deeper ones when the parameters are larger due to inherently non-adiabatic effects. Lastly, these phenomena, along with small parameter approximations [20] are shown to explain a general qualitative feature in performance [22] for deep QAOA circuits with slowly varying parameters.

The real-world limitations of NISQ processors, (e.g., the number of qubits, native gate sets, limited connectivity, effects of noise) have led the quantum computing community to design hardware-efficient ansätze, an example of hardware-algorithm co-design. In Ref. [23], a new ansatz that combines the mixing and phase-separation operators into a more general two-parameter family of operators is introduced, which is referred to as “Quantum Alternate Mixer-Phaser Ansatz” (QAMPA). The primary motivation for this new ansatz is to reduce the depth of circuits, which should improve the performance on NISQ processors. For fully connected binary quadratic optimization problems, the circuits compile to roughly half the depth of standard QAOA on quantum processing units (QPUs) with nearest-neighbor connectivity. Also, our numerical noiseless simulations demonstrate that QAMPA performs almost on par with standard QAOA in parameter regimes that are achievable in current hardware.

In Ref. [24], a Time-Block ansatz is proposed that requires only linear connectivity between qubits. Over-parametrization allows for more refined control over the circuits’ sublayers. As an additional contribution, the ordering of gates is used as a variational parameter, which is shown to greatly improve the performance in practice. Reported experimental implementation by the QuAIL team in collaboration with Rigetti Computings on Rigetti’s QPU pushes the state-of-the-art to a new scale, demonstrating variational quantum-classical optimization on 50-qubit systems using up to  $\sim 5,000$  two-qubit gates to approximately solve the fully-connected Sherrington-Kirpatrick model. Despite such complexities, the algorithmic performance gains are maintained over a random guessing solver given the same number of function calls. Furthermore, an increase in performance with circuit depth is observed.

In Ref. [25], the QuAIL team developed a novel meta-algorithm, Noise-Directed Adaptive Remapping (NDAR) for improved quantum optimization in the presence of certain types of hardware noise. The NDAR involves a feedback loop, where each step is a stochastic optimization, such as QAOA implemented on QPU. The algorithm assumes a known “attractor” state to which the noisy dynamics of a quantum device converges with time and noise accumulation. This is fulfilled experimentally in the case of, for example, amplitude damping, a type of noise prevalent in superconducting architectures. The aim of NDAR is to iteratively improve alignment between the low-energy states of the problem Hamiltonian and the attractor state of the noise. This is achieved by using the best solution from the previous step to re-map the problem Hamiltonian solved in the next step using so-called bitflip (or spin-reversal) gauge transformation. The obtained problem is equivalent to the original one, but the attractor state of the noise now becomes a better approximation to the ground state of the cost Hamiltonian. In cited Ref. [25], the team benchmarked NDAR experimentally on the newest generation of Rigetti’s superconducting chips, Ankaa-2. The reported approximation ratios are within the range 0.9-0.96 for random, fully connected graphs on  $n = 82$  qubits, using only depth  $p = 1$  QAOA applied in conjunction with NDAR. This compares to 0.34-0.51 for vanilla  $p = 1$  QAOA with the same number of function calls. As such, the obtained results are among the most complex and best-performant quantum optimization experiments performed to date (see, for example, [10] for an overview of recent experimental progress).

Applying QAOA algorithms to problems with constraints presents an implementation challenge for near-term quantum resources. Ref. [12] presented a general framework for designing constraint-specific mixers. In Ref. [26], strategies for enforcing hard constraints by using XY-Hamiltonians as mixing operators (mixers) are explored. Under the constraint that an integer variable admitting  $\kappa$  discrete values, encoded into qubits through the one-hot-encoding, XY Hamiltonian that preserves the total Z is a natural and efficient mixer choice. Despite the complexity of simulating the XY model, it is demonstrated that certain classes of the mixer Hamiltonian can be implemented without Trotter error in depth  $O(\kappa)$ . General strategies for implementing QAOA circuits on all-to-all connected hard-



ware graphs and linearly connected hardware graphs inspired by fermionic simulation techniques are also specified. The algorithmic performance is validated on graph coloring problems that are known to be challenging for a given classical algorithm. The general strategy of using XY-mixers is borne out numerically, demonstrating a significant improvement over the general X-mixer. The generalized W-state, an eigenstate of the XY mixer, which also corresponds to the symmetric (to particle permutation) state in the feasible subspace, yields better performance than easier-to-generate classical initial states when XY mixers are used. The XY mixer is further examined for graph coloring problem in the presence of noise in Ref. [27]. The evolution, which in the noiseless case would preserve the subspace defined by the symmetry, could break the symmetry and introduce non-valid states. The probability of staying in such symmetry-preserved subspaces under noise was analyzed, providing an exact formula for local depolarizing noise. These findings are applied to benchmark, under depolarizing noise, the symmetry robustness of XY-QAOA, which has local particle number conserving symmetries. The influence of the choice of the problem encoding on the symmetry robustness was also analyzed to provide guidance on realistic considerations in implementing symmetry-specific mixers for QAOA applications.

## 2.2. Entanglement verification in QAOA

A key ingredient of quantum algorithms is entanglement, although much about entanglement in quantum algorithms remains unknown. Verifying entanglement can be exponentially expensive. When information about the states to be verified is known, it may be possible to identify entanglement witnesses, operators that have bounds on their expectation values for all non-entangled states, and have expectation values that violate those bounds in the state whose entanglement is to be verified. A famous example is the observables involved in Bell inequalities. In the case of states arising in QAOA applied to MaxCut, Ref. [28] constructed entanglement witnesses that have only twice, or at most thrice, the number of terms as the underlying problem Hamiltonian. Moreover, in order to verify entanglement, one needs only measure QAOA-MaxCut states in 3 bases, measuring all qubits in the  $X$ ,  $Y$ , or  $Z$  basis. The work included experiments on Rigetti quantum processors that demonstrate this approach for verifying entanglement. The same work also developed signatures of coherence, another key property of quantum behavior, and demonstrated that the states obtained by running QAOA for MaxCut on Rigetti QPUs maintain coherence for sizable problems of interest.

## 2.3. Iterative Quantum Algorithms

Iterative quantum algorithms were originally proposed by Bravyi et al. [18] in the form of recursive QAOA (RQAOA), an extension based on QAOA. This algorithm solves MaxCut problems by using a quantum algorithm to identify the most probable edge to fix before reducing the problem size by fixing that edge. Our work extended this into a full framework of iterative quantum algorithms where the algorithm can be broken down into three parts [29, 30]:

1. Preparation Rule: prepare some information about the system. This can be preparing a quantum state using an optimization algorithm, sampling using a classical Monte Carlo, calculating graph properties, or anything else to gather information.
2. Selection Rule: using the information from the previous step, rank features of the problem, selecting a feature that you are most confident about altering or fixing.
3. Reduction Rule: Take the given feature and eliminate it from the problem. Usually, this involves fixing some information, such as collapsing two graph nodes into one. This reduction should be tracked and can be back-tracked through at the end of the procedure.

After each round of this procedure, the problem will be smaller and smaller until the optimal solution for this smaller problem is tractable to find. Then, this solution can be taken to back-track through the reduction rules to recover a candidate solution for the full problem.

In Ref. [29], these techniques are applied to a Maximum Independent Set problem, developing several variants of the Iterative Quantum Algorithm for this problem. Most notably, path-sum analysis techniques are applied to prove that an Iterative Quantum Algorithm using depth  $p = 1$  QAOA as its preparation rule only considers the same graph properties (i.e. degrees of all the nodes) as a classical greedy algorithm. Numerical simulations further verify that this quantum algorithm always makes the same selection choices that a classical greedy algorithm would, giving it the same performance and provable bounds as the classical algorithm. This is significant because it puts lower bounds and classical guarantees on the behavior of the quantum algorithm. At higher depths,  $p > 1$ , the classical side no longer offers such guarantees, but our numerical results show that the quantum algorithms outperform the classical greedy algorithm.

In [30], the QuAIL team worked with our hardware collaborator Rigetti Computing to introduce an iterative quantum heuristic optimization algorithm to solve combinatorial optimization problems. The quantum algorithm is implemented on Rigetti’s programmable superconducting quantum system using up to 72 qubits for solving paradigmatic Sherrington-Kirkpatrick Ising spin glass problems and observing an absolute performance better than a random sampling baseline.

Together with the Rigetti Computing team, we collaborated in [31] to propose a hybrid quantum-classical algorithm to solve the MaxCut problem on 3-regular graphs up to several thousand variables. Inspired by semidefinite programming and considering the structure of the problem, the proposed method solves problems beyond the number of qubits in Rigetti’s quantum computer. This approach was compared to state-of-the-art classical methods and achieved an average performance of 99% on a large set of instances, making it competitive with respect to high-performance classical methods.

## 2.4. Quantum Annealing Advances

Quantum annealing has been a major research focus area of the QuAIL team since its earliest days [2, 3]. Quantum

annealing [32, 33] is a general and universal [34] setting for quantum computation that relies on continuous time evolution of a quantum system under mixing and problem Hamiltonians. Quantum annealing is similar to QAOA discussed before but relies on slow evolutions and quantum adiabaticity. This evolution can often be realized on physical devices that require embedding of the problem into the native structure of the device, and often fine tuning of the annealing schedule is needed to realize improvements. Quantum annealing has so far demonstrated quantum speedups over classical optimization algorithms only in very limited circumstances. One reason is the need for minor embedding which maps optimization problems of interest (e.g. SAT, vertex-cover) to the type of problem natively optimized by a quantum annealer, an Ising Hamiltonian with connectivity determined by the architecture, which usually is an Ising spin Hamiltonian of the form

$$H = - \sum_{\langle i,j \rangle \in E} J_{ij} \sigma_i^z \sigma_j^z - \sum_{i \in V} h_i \sigma_i^z. \quad (1)$$

The qubit operators  $\sigma_i^z$  label the vertices of a graph  $G = (V, E)$  and the two-qubit interaction terms  $J_{ij} \sigma_i^z \sigma_j^z$  are associated to the edges  $\langle i, j \rangle \in E$  of  $G$ . While the coupling constants,  $J_{ij}$ , and the local fields,  $h_i$ , can be tuned to a degree, the graph  $G$  itself is fixed once and for all by the architecture of the machine and is limited by engineering considerations. If the target cost function of the problem of interest is not of the form in Eq. (1), one needs to devise a way of mapping it to a cost function that the machine can accept. This usually requires using auxiliary qubits and representing the logical values of the original problem as long chains of physical qubits coupled together. This mapping to a more complicated system was observed to significantly reduce the performance of quantum annealers compared to their performance on their native problem due to the freezing of such chain in the “wrong” configurations, in a way that the quantum annealer is no longer able to change their states dynamically.

Many practical optimization problems have high connectivity. Minor embedding addresses this mismatch by effectively increasing the connectivity by grouping qubits together to act as one logical qubit, and adding terms to the cost function that penalizes configurations in which the qubits making up a logical qubit do not align. Parameters are chosen so that the ground state is preserved by the embedding.

In [35], certain practical aspects of embedding for optimization problems through experimental demonstrations on a commercial quantum annealer were investigated. The effects of embedding and its interplay with other annealing parameters (particularly those related to the annealing schedule) on the performance of the device for solving optimization problems were studied, aiming to provide both deeper insights into the physics of quantum annealing devices as well as pragmatic recommendations for their use. Two performance-enhancing methods previously applied to native optimization problems were adapted and shown to be advantageous also for embedded ones. A novel approach to enable gauge transformations for problems with the qubit coupling strength  $J$  in an asymmetric range is introduced, making it amenable to embedded problems. An anneal-

ing schedule with an appropriately located pause was also confirmed to still improve performance in the embedded case and explored how the optimal location shifts with the magnitude of the ferromagnetic coupling  $|J_F|$ , thus extending the theoretical picture for pausing and thermalization in quantum annealing to the embedded case.

Our follow-up work [36] delved deeper into this theory and practical recommendations through demonstrations of an updated annealing architecture and several problem classes requiring embedding. Various aspects of the physics-based picture previously explored were confirmed, and taken further by identifying certain characteristics of an optimization problem that are predictive of its hardness for currently available quantum annealing devices. Based on these results, a set of qualitative guidelines for parameter setting in quantum annealers was presented that is expected to be particularly useful for users coming into quantum annealing from other areas of optimization, and that should help maximize the performance of these devices without requiring a prohibitively large amount of resources.

Also focusing on practical aspects and existing annealing devices is [37], in which the performance of quantum annealers has improved since their inception was explored, with a comparative study carried out throughout the years using each new available processor. This work highlights how the improvements in connectivity, leading to smaller embeddings, boost the optimization capabilities of these machines.

While the embedding process for optimization is sound by construction, the same is not true for sampling. Ref. [38] showed that even with an annealer that returns a perfectly classical Boltzmann (thermal) distribution of the final Hamiltonian  $e^{-H/T}$ , the process of embedding drastically alters the statistics. Not only is the distribution ‘hotter’ (larger  $T$ ), it is also strictly not Boltzmann anymore. A follow-up work [39] applied a similar analysis of the quantum version of the problem (sampling a quantum Boltzmann distribution, where the Hamiltonian has non-zero off-diagonals), which showed similar effects; quantum observables were distorted by the embedding, including shifting the location of a phase transition. The message in these works is clear: annealers with naive embedding can not be used for classical or quantum thermal sampling.

While these sampling capabilities are still limited, they can be useful for certain scenarios. For native problems (i.e. not needing embedding), it was demonstrated in [40] that using a pause and quench schedule, can be used to probe the thermal properties of the strictly quantum Hamiltonian that describes the system at the time of the pause, allowing us to distinguish between graphs whose classical Ising spectra are the same, which could not be done with a standard annealing schedule. The theoretical work [41] demonstrates that embedding chains typically crop up as harmful Griffiths phases but can also be used as a resource to balance out singularities in the logical problem, changing its universality class.

Ref [42] used constrained programming to find provably optimal embeddings for quantum annealers. A mixed-integer programming formulation was also developed, and a comparison of these approaches was performed to find embeddings of

various problems in D-Wave’s quantum annealers. Although these approaches are generally not capable of finding embeddings as efficiently as the heuristics, these methods were well suited for those cases that were pathologically difficult for the heuristic methods,

In quantum annealing, finding a good schedule is often essential to obtaining quantum advantage. Like finding good parameters in QAOA, optimizing the schedule can be itself a difficult task. Therefore, it is essential to create ansätze that can be effectively optimized with as few variational parameters as possible. QAOA serves as an effective ansatz in a gate-based computer, but a quantum annealer has additional degrees of freedom in its schedule. Ref. [[43]] introduced a new parameterization of the quantum annealing schedule based on clipped polynomials.

In addition to introducing this schedule parameterization, a benchmarking against QAOA was performed, introducing notions of emulation time. Additionally, the ability of this polynomial schedule to numerically emulate QAOA was shown, producing an equivalent or better final state using the same number of variational parameters. Further, the polynomial schedule can produce states that are strictly better than those produced by QAOA. This shows increased power for annealing schedules, but it requires much more experimental and physical control.

Ref. [[44]] derived a version of quantum speed limits specific to quantum annealing. The work explored this bound in several toy problems and showed that this bound is saturable for some quantum annealing procedures, including the Hamming spike, the  $p$ -spin model, and the adiabatic unstructured search. These problems all represent extreme toy problems, so further refinements are needed to make this bound saturable for more realistic and less symmetric problems.

Further advances related to quantum annealing are discussed in Secs. 6.4 and 6.5.

## 2.5. Quantum Optimization Applications

The QuAIL group has a long history of working on planning-related problems [[2, 3]]. Extending our previous work for planning problems, [[45]] presents novel efficient QAOA constructions for optimization problems over proper colorings of chordal graphs. The effectiveness of these constructions was demonstrated using the flight-gate assignment problem, in which flights are assigned to airport gates to minimize the total transit time of all passengers. Feasible assignments correspond to proper graph coloring of a conflict graph derived instance-wise from the input data. Further, [[46]] investigated the feasibility of applying quantum annealing to solve a simplified air traffic management problem (strategic conflict resolution) for wind-optimal trajectories. Our mapping is performed through an original presentation of the conflict-resolution problem in terms of a conflict graph, where the nodes of the graph represent flights and the edges represent a potential conflict between flights.

In a pair of papers, [[47, 48]], we were able to speed up certain highly utilized subroutines for solving constraint satis-

faction problems, given access to a fault-tolerant quantum computer. One key result leveraged quantum search to accelerate filtering for the ‘alldifferent’ constraint (all variables take on unique values). This work showed how to incorporate quantum filtering algorithms into a hybrid classical-quantum backtracking search protocol. Resource estimates suggest constraint programming is a promising candidate application for early fault-tolerant quantum computers.

In [[49]], the relationship between QAOA and the underlying symmetries of the objective function is explored, particularly in hard problem classes where a nontrivial symmetry subgroup can be obtained efficiently. The work shows how symmetries of the objective function imply invariant measurement outcome probabilities across states connected by such symmetries, independent of the choice of algorithm parameters or number of layers. Using this result, machine learning techniques can predict QAOA performance based on symmetry considerations. Numerical evidence suggests that a small set of graph symmetry properties suffices to predict the minimum QAOA depth required to achieve a target approximation ratio on the MaxCut problem in a practically relevant setting where QAOA parameter schedules are constrained to be linear.

## 2.6. Distributed Quantum Computing

We have also looked at distributed graph problems, in which a network of computers, each with some information about the graph, collaboratively work to compute a global graph property. Ref. [[50]] presents two algorithms in the Quantum CONGEST-CLIQUE model of distributed computation that succeed with high probability: one for producing an approximately optimal Steiner tree and one for producing an exact spanning arborescence of minimum weight, the analog of a minimum spacing tree in a directed graph.

The CONGEST distributed computational model allows messages of limited size to be transmitted within a network described by a communication graph of size  $n$  in a series of rounds to address a computational problem. The size limitation for such messages is  $O(\log(n))$  bits at each edge of the communication graph per round. The communication graph in the CONGEST-CLIQUE model is fully connected. In the quantum CONGEST-CLIQUE model, at most  $O(\log n)$ , classical and quantum bits (qubits) can be communicated across each edge of the communication graph per round.

Each of these algorithms uses  $O(n^{1/4} \log^k n) = \tilde{O}(n^{1/4})$  rounds of communication and  $\tilde{O}(n^{9/4})$  messages, achieving a lower round and message complexity than any known algorithm in the classical CONGEST-CLIQUE model. At a high level, these results are achieved by combining classical algorithms with fast quantum subroutines. These asymptotic speedups further contribute to understanding what problems can be solved more efficiently when quantum communication is allowed in this CONGEST-CLIQUE model of distributed computation.

As a final result, the pre-factors accompanying the poly-logarithmic terms derived in the complexity analysis were estimated. By performing such an analysis, it was discovered that, although providing an asymptotic speedup, these distributed



quantum methods would only surpass the performance of trivial classical algorithms when the graphs to be processed (i.e., the network of quantum computers) would have over  $10^{21}$  nodes. Such results only motivated us to work on improving these distributed algorithms, for which our current work addresses generalizations of these methods.

### 2.7. Quantum Algorithm Design and Problem Encoding

QAOA and other quantum approaches to optimization can also be applied to problems on  $d$ -level variables. Building off of previous work that considered the qubit case in detail [[12, 51]], Ref [[52]] presents a toolkit for design and analysis of discrete (i.e., integer-based) optimization problems, wherein the problem and corresponding algorithmic primitives are expressed using a quantum intermediate representation that is encoding-independent. This compact representation often facilitates efficient problem compilation and comparison between different encoding choices, for example between variants of binary and unary encodings for fixed  $d$ . Numerical studies comparing several qubit encodings exhibit a number of preliminary trends toward guiding the choice of encoding for a particular set of hardware, problem, and algorithm. For moderate sized up to 16-level quantum variables, low-depth mixing operators are constructed, demonstrating that binary encodings remain amenable for QAOA.

Diagrammatic representations of quantum circuits offer novel approaches to their design and analysis. Extensions of the ZX-calculus [53] especially suitable for parameterized quantum circuits are proposed in [[54]], in particular for computing observable expectation values as functions of the quantum circuit parameters, which are important algorithmic quantities in a variety of applications ranging from combinatorial optimization to quantum chemistry. In particular, formal rules for dealing with linear combinations of ZX-diagrams are given, where the relative complex-valued scale factors of each diagram must be kept track of, in contrast to most previously applications. The diagrammatic approach is shown to offer useful insights into algorithm structure and performance through direct application to several example applications drawn from the literature including realizations of hardware-efficient ansatzes and QAOA, where calculations can become more intuitive and potentially easier to approach systematically than by alternative means. Recent related work has further considered applying similar diagrammatic approaches to the study of barren plateaus [55] and to quantum machine learning [56].

### 2.8. Physics Inspired Heuristics

One of the first and most significant boons of quantum computing has been its inspiration and advancement of classical computing methods. Physics-inspired heuristics provide completely classical methods of solving hard problems in new ways. As part of our work in this area, a set of tools and methods that not only implement some of these physics- and quantum-inspired heuristics, but also integrate them with other solution techniques to address challenging computational

tasks, using quantum computing only as a source of inspiration have been developed.

As mentioned above, there is a significant effort from the QuAIL team to understand quantum algorithms for optimization, namely quantum annealing and QAOA. These methods are designed to find ground states of transverse field Ising problems, which, from an optimization standpoint, correspond to QUBO problems. A classical algorithm to address these problems is parallel tempering, a modification of simulated annealing that has proved to be one of the strongest contenders when comparing the performance of quantum methods for optimization. An efficient Python implementation of the parallel tempering algorithm has been provided as part of the open-source project PySA [[57]].

Another contribution is in the tools required to address constrained optimization problems using physics- and quantum-inspired heuristic methods. Although there are known mappings of constrained optimization problems to the QUBO formalism, which is amenable for these heuristics and other quantum algorithms, their efficient implementation is challenging. The group has also developed tools for the community trying to use and understand these methods when facing applications, usually represented through constrained optimization. These tools are provided as software written in the Julia programming language, under the umbrella of the package QUBO.jl [[58]]. Within it, there are tools for reformulating mixed-integer nonlinear programming problems into QUBO, connecting these instances with classical and quantum solvers, and analyzing both the instances and results provided by the solvers.

Although mapping constraints as penalizations might be sufficient for their enforcement, this method requires finding these penalization factors, which might lead to challenging numerical issues for both classical and quantum methods. A way of leveraging these Ising solvers, potentially quantum, while enforcing the constraints exactly for mixed-binary quadratic programs was proposed. This method relies on a copositive reformulation of the problem and the application of a cutting-plane algorithm, which generates increasingly accurate linearizations of the problem, for which center solutions are computed classically, derived from the possibly suboptimal solutions of QUBO problems found by the Ising solvers. This method, described in [[59]], also provides convergence guarantees to provable optimal solutions despite the heuristic nature of the Ising solvers, and for a test case of finding the Max-clique of random graphs, it is able to match and even surpass state-of-the-art solvers for these constrained problems.

On the other hand, the group has also worked on the simulation of coherent Ising machines (CIM), a non-conventional architecture for solving Ising problems heuristically. The CIM dynamics are described by a set of ordinary differential equations, which can also address continuous non-convex quadratic optimization. The dynamics of such continuous variable CIMs (CV-CIM) use optical pulses to perform the optimization. Updates to those pulses are computed via stochastic gradient descent. Classical optimization techniques have been used to improve the dynamic simulations of CV-CIM [[60]], namely momentum and Adam updates, to address the weaknesses of



stochastic gradient descent found in classical optimization. Through this modification, the CV-CIM's convergence, sample diversity, and stability can be significantly improved.

Continuing with the CIM work, the physics and the performance of this class of machines, which are faithfully described by a set of coupled ordinary differential equations ODEs, was investigated as part of the NTT Phi Lab collaboration and the NTT NSF Expeditions program on Coherent Ising Machines led by Stanford University. The fast clock speed and relatively mature engineering in the field of photonics allows for the practical possibility of these systems, either simulated in CPUs, GPUs, FPGAs or realized via integrated photonic circuits to be deployed as components to solve contrived problems that have crucial time-threshold, such as the MIMO wireless Maximum-Likelihood decoding problem already subject of Ref. [[61]] (more on this below). Indeed, elaborated fully-electronic or opto-electronic implementation of the CIM machines are at the moment among the most performant Ising solvers known for a variety of challenging benchmarks, according to a recent review [[62]]. In [[63]], the team investigate the basic numerical implementation of the CIM equations for the purpose of wireless decoding, setting up problems using the Rayleigh fading channel model and calculating resource requirement for a high-speed Ising machine to increase the overall throughput of current state-of-art network implementations. It is noteworthy that the work spurred multiple derivative improvements from other groups, and led to current related lines of research active, such as a full-FPGA optimized implementation and a neural-network representation of the CIM using Deep Operator Networks [[64]].

The sign problem is one of the biggest impediments to simulating quantum systems using Monte Carlo techniques. When Monte Carlo techniques are used to simulate a Hamiltonian or system with a sign problem, the pseudo-probabilities for the Monte Carlo method oscillate rapidly between positive and negative values (or arbitrary complex phases when simulating real-time evolution), making evaluation of averages and integrals exponentially difficult. Most Monte Carlo techniques for simulating quantum systems are, therefore, limited to sign-problem-free Hamiltonians, a vanishingly small fraction of all possible systems.

There are methods for mitigating but not solving the sign problem, such as one already used in high energy physics, involving integration along Lefschetz Thimbles [65, 66]. These thimbles are a generalizations of the complex plane of stationary phase. In broad terms, these methods seek to evaluate a highly oscillatory integral with a sign problem by deforming the surface of integration into the complex hyper-plane and deforming the surface iteratively to try to get it to flow to a Lefschetz Thimble. The idea then is that on this thimble, the integral is easy to evaluate since its phase is no longer oscillatory. These methods are themselves highly computationally taxing, and at best they can mitigate the slow-down of the sign problem to a lesser exponential factor. Our work with collaborators [[67]], sought to extend these ideas from continuous variable high energy physics to the discrete variables of quantum spin systems. This process is non-trivial since the complexification requires

continuous variables, requiring us to look at over-parameterized continuous bases for quantum spin systems, such as the spin coherent states, which themselves require approximations that are sometimes ill-justified for low spin number. Our work successfully implemented these Lefschetz Monte Carlo methods and demonstrated their effectiveness for a small three-spin system with a sign problem.

An example of physics-inspired heuristic applied to real-world application is the use of the parallel tempering algorithm for soft MIMO decoding in 5G (ParaMax) [[61]]. Here, we demonstrate that ParaMax can achieve near optimal maximum-likelihood throughput performance in the Large MIMO regime, Massive MIMO systems where the base station has additional RF chains, to approach the number of base station antennas, in order to support even more parallel spatial streams. These results are achieved by using the ParaMax Ising Solver (PMIS). It is based on simulated annealing, featuring a parallel tempering algorithm highly-tailored to optimize the Ising model of MIMO detection. While the front-end of our PMIS implementation is in Python, the core is completely written in C++. To further maximize the performance of PMIS to satisfy limited processing time in wireless standards, the following innovations have been implemented: the use of static memory to improve compile-time optimization; the use of parameter pack expansion to unroll vector-vector and matrix-vector multiplications; and the use of SIMD instructions to further improve operations like vector-vector and matrix-vector multiplications.

One of the main reasons classical heuristics fail to identify solutions of optimization problems is being stuck in local minima [[68]]. Thermal cycling is an optimization algorithm based on simulated annealing that uses cycles of heating and quenching, while following a decreasing temperature schedule, to escape local minima and improve the performance of local heuristics [[69]].

In [[70]], a comprehensive parameter tuning of the algorithm was performed and it was demonstrated that it competes closely with other state-of-the-art algorithms such as parallel tempering with isoenergetic cluster moves [[71, 72, 73]], while overwhelmingly outperforming more simplistic heuristics such as simulated annealing. In details, the thermal cycling algorithm works by preparing an ensemble  $\Omega$  of  $N_p$  lower energy states among  $N_0$  quenched random configurations. Starting from the initial inverse temperature  $\beta_i = 0$ , the above pool of states is annealed toward a final inverse temperature of  $\beta_f$  in  $N_T$  steps. At any given temperature step, a given number of states in  $\Omega$  are update using Metropolis (heating) followed by an immediate quench via a local search method (cooling). If any of the new states have an energy smaller than its original state, such states are updated in  $\Omega$ . In practice, the above process steers the ensemble toward the low-lying states while ensuring that meta-stable configurations do not hinder the dynamics. The temperature is then reduced, and the cycles start over.

The main advantage of algorithms like the thermal cycling is the possibility of using a variety of variable-update classes in the quenching phase. By using carefully designed updates, the exploration of the exponentially many meta-stable can be

avoided in favor of lower-energy states. However, considering the extra computational cost associated with tailored updates, it is important to considering the trade-off between a better exploration of the configuration landscape and the induced overhead. By accurately choosing the optimal update strategy, it was demonstrated that the thermal cycling algorithm is a competitive heuristic when complex structures are present, avoiding being trapped in local minima. While the thermal cycling algorithm performs the worst when narrow barriers are present, where algorithms like simulated quantum annealing and isenergetic cluster updates perform the best, it was observed that the thermal cycling algorithm works the best on dense graphs with wide barriers.

### 3. Quantum-compatible Machine Learning Algorithms

Over the past few years, significant strides in the field of machine learning (ML) have been made, particularly in integrating quantum computing with ML algorithms (QML). This fusion has been explored in both discriminative models like classification and generative models, focusing on understanding and leveraging quantum probability distributions. The central challenge has been effectively integrating samples from quantum processors into state-of-the-art ML models, such as Variational AutoEncoders (VAE) and Generative Adversarial Networks (GAN). The QuAIL team has worked with collaborators, from domain experts to hardware groups, to advance these algorithms. These efforts on specific applications in aeronautics and Earth science including image segmentation and modeling for wildfire detection, time-series analysis for flight anomaly detection, and parameter learning for material science modeling. In all cases, there is a tight connection between the algorithm and the hardware, illustrating the potential for specialized quantum hardware that have less stringent requirements than universal quantum processors and so could be built more easily and put to use earlier than universal processors, and also the potential of algorithms that can be tailored to emerging quantum hardware. This connection enables a virtuous cycle of hardware-algorithm co-design.

This section delves into these hybrid ML models, examining how quantum samples, particularly from Boltzmann Machines, are integrated within the latent spaces of these models. In several instances, these quantum-compatible methods, based on discrete latent space models, have provided significant value over standard approaches, demonstrating continued opportunity in QML research. However, the utilization of samples from noisy, sparse latent distributions, characteristic of modern quantum hardware, remains a challenge. Overcoming these hurdles is vital, especially considering the demands of high-performance computing environments.

There are four distinct paradigms for how to combine machine learning with quantum computing:

1. **Classical-Classical (CC):** This approach leverages quantum-inspired methods within discrete latent space models, showcasing how quantum influence can enhance traditional algorithms.
2. **Classical-Quantum (CQ):** Exemplified by models such as the Quantum-assisted Variational Autoencoder (QVAE), this paradigm involves the use of quantum computing elements to process and interpret classical data.
3. **Quantum-Classical (QC):** Models optimizing quantum heuristics with classical machine learning techniques and reflects the application of classical ML techniques to optimize or interpret results obtained from quantum simulations or quantum processors. This paradigm reflects the integration of quantum computational power in solving complex quantum problems, which are then analyzed and optimized using classical ML techniques.
4. **Quantum-Quantum (QQ):** Although not explored in this section, this approach signifies the complete integration of quantum data with quantum computational methods.

With the exception of subsection 3.5, this section mostly concentrates on advancements in CQ paradigm.

#### 3.1. Quantum-Assisted Variational Autoencoder

Both quantum and classical Boltzmann Machine (BM) are able to model powerful and flexible probability distributions. However, training these models classically is challenging and time-consuming. Quantum annealers, regarded as analog quantum devices, can effectively simulate Quantum BMs (QBMs). Ref. [[74]] introduced a quantum-assisted VAE (QVAE), which incorporates samples from a quantum annealer into the prior of a discrete VAE. The QBM component of the model features a tunable transverse field, enhancing the prior's expressiveness and allowing for a more focused latent distribution without compromising the network's learning ability. This model excels at handling high-dimensional data spaces, a significant challenge for classical machine learning algorithms [[74]].

The QVAE's encoder comprised two fully-connected layers leading into a hierarchical posterior, facilitating complex posterior distributions and a tighter Evidence lower Bound (ELBO). To prevent overfitting, the decoder has a simpler structure with two fully-connected layers [[74]].

An integral part of the study was comparing the performance of a Restricted BM (RBM), a classically simulated QBM, and a QBM trained with samples drawn from the D-Wave 2000Q quantum annealer. This device, featuring 2048 qubits each of degree 6, required embedding any RBM larger than 6 nodes per side with a coupling strength of -1. The logical value of a variable at the end of an anneal was determined by a majority vote, adding a layer of robustness to the modeling process [[74]].

After training, the QVAE's encoder indexed the dataset efficiently without extra discretization due to its binary latent space. The encoding algorithm clustered similar objects under identical bit strings and an inverted index was constructed to map each bit string to its corresponding data points. Since each item is processed once and independently, the original data can be removed from the main memory, reducing memory usage.

For queries, the QVAE encodes the item and sorts occupied bit strings by their Hamming distance to the query embedding, balancing memory efficiency and search resolution. The results were significant [[74]]:

- **Embedded Proximity and Hamming Distance:** Experiments validated that Hamming distance in the compressed space effectively approximates Euclidean distance in the original space, as shown through k-Approximate Nearest Neighbor Search (ANNS) on the Moderate Resolution Imaging Spectroradiometer (MODIS) dataset. This underscores the efficiency of quantum-assisted techniques in high-dimensional data retrieval.
- **Impact of the Transverse Field:** Adjusting the transverse field parameter influenced the distribution’s characteristics and search speed, with optimal speedup observed at certain ranges. This highlights the precise control quantum models offer in data representation and optimization.
- **Memory Consumption:** The quantum model showed superior memory efficiency over methods like Hierarchical Navigable Small World (HNSW) and Locality Sensitive Hashing (LSH), particularly for large datasets like the complete MODIS dataset, underscoring the models practicality and scalability.

The effective utilization of quantum annealing in handling latent spaces by the QVAE indicates potential quantum enhancements for other ML models, like the supervised U-Net-VAE hybrid used in Ref. [[75]] for wildfire detection. This model, integrating a U-Net architecture with a VAE, comprises four submodels: (1) a prior network, (2) a posterior network, (3) a U-Net network for feature extraction, and (4) a combination network. It is effective in generating stochastic wildfire segmentations and simulating unknown wildfire scenarios. Incorporating quantum-sourced samples could further enhance its predictive accuracy and robustness, particularly in analyzing high-dimensional satellite imagery, offering a promising direction for future research in remote sensing and wildfire analysis [[75]].

### 3.2. Optimizing Adversarial Networks Through Quantum Annealing

In exploring the integration of quantum annealing with deep learning architectures, Ref. [[76]] introduced the Quantum-assisted Associative Adversarial Network (QAAAN). This model incorporates quantum annealing to train a BM that optimizes the feature distribution extracted by the discriminator network in the adversarial model. The unique capabilities of quantum computing are leveraged to find more effective representations in the latent space, thereby enhancing the generation of realistic data. This approach, particularly significant for its exploration of reparametrization of discrete variables, demonstrates a crucial step in integrating quantum models with traditional neural networks [[76]].

In the QAAAN study, the researchers investigated various topologies for the probabilistic graphical models in the latent space, including complete, symmetric bipartite, and Chimera topologies. It is noteworthy that the choice of topology affects both the model’s learning rate and quality, similar to how transverse fields influenced the QVAE’s performance in Ref. [[77]]. This insight is crucial for future developments in quantum-enhanced deep learning, as it highlights the need for careful consideration of model architecture to harness the full potential of quantum computing.

The QAAAN’s exploration of topological variations extends to practical performance metrics. The researchers utilized the Inception Score and the Frechet Inception Distance to assess the model’s ability to generate realistic data. This evaluation was pivotal in establishing the QAAAN’s efficacy in real-world applications. The research also demonstrated the model’s scalability by successfully applying it to complex datasets, such as the LSUN bedrooms dataset. This scalability indicates the potential for quantum-assisted models in handling large-scale, high-dimensional data challenges [[76]].

### 3.3. Advancements in Quantum-Inspired Generative Models: RBMs in VAEs and Invertible Flows

Researchers in [[78]] introduced models that mark a significant advancement in combining Energy Based Models (EBMs) with Invertible Flows (IFs), specifically RBM-Flow and D-Flow. These non-autoregressive IFs enable exact likelihood training and efficient sampling without requiring a discriminator network, which is a notable difference from VAEs, autoregressive models, and GANs, respectively [[78]].

RBM-Flow, a subclass of EBM-Flow, leverages RBM as its trainable, base distribution, thus enhancing the model’s capacity to capture complex data distributions. D-Flow, an evolution of RBM-Flow, zeroes all couplings in the latent RBM, allowing the encoding of global features as discrete labels in the latent space, offering a structured approach for meaningfully handling high-dimensional data. The usage of RBMs in RBM-Flow primarily enhances their energy-based properties and discrete nature to boost the generative capabilities of the IF framework [[78]].

### 3.4. Quantum-Compatible Discrete Deep Generative Models

In [[79]], deep generative learning models are developed and applied to the task of anomaly detection in a commercial flight-operations dataset consisting of multivariate time series. Specifically, the performance of three unsupervised deep generative models are explored which consist of variational autoencoders with Gaussian, Bernoulli, and Boltzmann priors. Two of the VAEs contained discrete latent variables (DVAEs), one with a factorized Bernoulli prior and one with a RBM prior. The work demonstrated the competitiveness of a discrete deep generative model with its Gaussian counterpart for anomaly detection tasks. Also of note, the DVAE model with RBM prior can be readily integrated with quantum sampling by replacing its generative process with measurements of quantum states obtained from quantum hardware devices, such as a quantum annealer or gate-model device.



### 3.5. Quantum-Enhanced Optimization of Quantum Heuristics

VQAs, categorized under quantum heuristics, show great potential for practical quantum computing applications. The optimization of these algorithms for effective hardware performance is a critical area of focus. Ref. [[77]] assessed the efficacy of a Long Short Term Memory (LSTM) recurrent neural network model (the meta-learner) in optimizing two quantum heuristics, comparing its performance traditional optimizers (Bayesian optimization, evolutionary strategies, L-BFGS-B and Nelder-Mead).

The meta-learner outperformed traditional optimizers, such as Bayesian optimization, evolutionary strategies, L-BFGS-B, and Nelder-Mead, demonstrating superior performance in noisy environments. For example, starting from the around the same point in Fermi-Hubbard model problems, the L-BFGS-B performance reduced by 0.35 whereas the meta-learner only reduced by 0.2 [[77]], indicating that meta-learning will be especially useful in noisy near-term quantum heuristics implemented on hardware. In addition to robustness, the meta-learner showed a higher frequency of reaching near-optimal solutions than the next best optimizer (evolutionary strategies) in noisy simulations [[77]].

Looking ahead, the continued improvement of meta-learning methods is anticipated. Despite the current lack of investigation into their performance scaling to larger problem sizes, largely due to the challenges in simulating large quantum systems, the potential for these methods on hardware implementations is significant. The meta-learner introduced by [[77]] applies a single model across various parameters (a 'coordinatewise' approach). Envisioning a 'qubitwise' approach, where distinct models are trained for each qubit's parameters in a given hardware graph, could open up new optimization avenues. Such a method might account for the unique physical characteristics of each qubit, potentially leading to more finely-tuned optimizations tailored to specific hardware environments. This direction underscores the importance of developing meta-learning methods that not only adapt to the complexities of quantum problems but also leverage the peculiarities of quantum hardware.

Ref. [[77]] also notes the different implementations of QAOA used for Graph Bisection and MAX-2-SAT. Understanding the impact of mixer and initial state variations on performance, as well as characterizing the relative power of different QAOA mixers, remains an open area of research.

## 4. Quantum Computing for Simulation

Our recent work related to quantum simulations for condensed matter physics, quantum computing for material discovery, quantum simulation for chemistry, and open quantum system simulation are highlighted in this section. Among the many advancements showcased in this section are recent works that use variational techniques for both weakly and strongly correlated systems that are able to reach chemical accuracy with a minimal number of time steps; extension of existing neural network techniques for solving electronic structure problems; and

a further reminder of the importance of classical HPC computing methods for benchmarking and quantum chemistry.

### 4.1. Quantum Simulation for Condensed Matter Physics

#### 4.1.1. Kitaev Spin Models

Quantum spin systems, notably the Kitaev honeycomb model, are promising candidates for demonstrating the potential capabilities of quantum computers over classical ones in applications beyond simulation. Kitaev's influential work [80] introduced this well-studied model, which is especially attractive for quantum algorithm research due to its simple lattice structure and the simplicity with which its Hamiltonian can be expressed in terms of Pauli operators.

In the study [[81]], the team successfully prepared the ground state of the Kitaev honeycomb model using the Hamiltonian variational ansatz (HVA) [82]. The model's square octagon lattice version was found to be particularly well-suited for implementation on Rigetti's Aspen-9 quantum processor. In practical experiments, their small-scale demonstrations of developed techniques, enhanced with error mitigation strategies, yielded results that closely aligned with theoretical predictions.

Further exploration into the Kitaev honeycomb spin model revealed an alternative approach [[83]], wherein the Hamiltonian is expressed purely in terms of fermionic degrees of freedom. This alternative fermionic description of the spin-system not only confirmed the model's exact solvability in the absence of interactions, but also led to the development of tailored variational quantum eigensolver (VQE) ansätze for different types of interactions. This approach allowed for accurate ground state preparations that aligned with theoretical results and demonstrated that for certain interactions, the number of required qubits could be cut in half [[83]].

#### 4.1.2. Fermi-Hubbard Model

The Fermi-Hubbard Model (FHM) is crucial in quantum physics for exploring phenomena like superconductivity and insulating phases in materials. It describes the complex interactions of electrons moving within a lattice, offering critical insights into these states.

A significant study [84] simulated the dynamics of a one-dimensional FHM using a 16-qubit digital superconducting quantum processor. The simulation revealed distinct spreading velocities for charge and spin densities in highly excited regimes, surpassing the limits of the conventional quasiparticle model. To mitigate systematic errors and decoherence, the researchers developed an accurate gate calibration procedure and employed a series of error-mitigation techniques. These procedures allowed for the faithful simulation of the model's time evolution, even with over 600 two-qubit gates in the circuits. This work represents a crucial step towards practical quantum simulation of strongly correlated phenomena using current quantum devices, such as the classically hard 2D FHM [84].

Subsequently, [[85]] explored a quantum annealing protocol for the FHM using a novel fermion-to-qubit encoding,



which simplifies the FHM representation and enhances simulation. The proposed driver solutions targets high fidelity in various FHM interaction regimes, especially in intermediate and large  $U/t$  regimes. The protocol is designed for coherent evolution fermionic logical subspace, particularly suited for accurate simulation [[85]].

The study discussed potential scalability and hardware noise challenges, noting the lack of currently suitable hardware. However, the research suggested that advancements in quantum annealing hardware could enable practical implementation, positioning this method as a promising contender to gate-based algorithms for quantum advantage in non-fault tolerant regimes [[85]].

## 4.2. Quantum Computing for Material Discovery

### 4.2.1. Advancements in State Preparation Techniques

Quantum computations offer significant advancements in simulating quantum many-body systems, with various approaches focusing on the efficient preparation of ground or excited states.

In the study [[86]], researchers explored a novel method for simulating the time dynamics of Adiabatic State Preparation (ASP). ASP, which is an alternative approach to preparing ground states to study their wave-function and energy, starts with a known, easily prepared ground state Hamiltonian and evolves it slowly to the target Hamiltonian, maintaining the system in its ground state, as per the quantum adiabatic theorem (ensuring the system remains in the ground state throughout the evolution). The approach employs an adaptive sampling configuration interaction scheme, enhancing the efficiency of preparing ground and excited states, and is particularly effective in avoiding the complexities of optimization in noisy environments.

In [[87]], a systematic investigation into VQEs for determining ground-state energies and properties of two-dimensional model fermionic systems was presented. The research focuses on the efficiency of different entangler blocks in the VQE and how they impact the convergence to the system's ground state with minimal gate operations, crucial for NISQ devices. The study also monitors the entanglement during optimization using the concurrence measure and investigates the scaling of the VQE circuit depth as a function of the desired energy accuracy [[87]]. In related work, [[88]] presented and benchmarked an approach for finding good starting parameters for parameterized quantum circuits by classically simulating VQE by approximating the parameterized quantum circuit as a matrix-product-state (MPS) with a limited bond dimension.

Lastly, in [[89]] a detailed analysis of variational quantum phase estimation (VQPE), a method rooted in real-time evolution for estimating ground and excited states on near-term quantum hardware. The study establishes the theoretical foundation of VQPE and highlights its efficiency in creating compact variational expansions for solving strongly correlated Hamiltonians. Central to VQPE is a set of equations with a geometrical interpretation, crucial for time evolution grid setup to decouple eigenstates from time-evolved expansion states. This

connects VQPE to classical filter diagonalization algorithms. The paper also introduces a unitary formulation of VQPE, reducing the required quantum measurements, and analyzes the impact of noise, significantly improving upon previous considerations. Through both numerical simulations and a hardware implementation for the transverse field Ising model, the paper demonstrates VQPE's efficiency for both weakly and strongly correlated systems, reaching chemical accuracy with a minimal number of time steps [[89]].

### 4.2.2. Many-body Scar Preparation

Many body scars are an interesting phenomenon that seemingly violate the eigenstate thermalization hypothesis, similar to the phenomenon of many-body localization. Such scarred states give rise to long-lived coherent dynamics in an otherwise thermalizing system and may provide a path towards physical realizations of quantum memory.

In Ref. [[90]], state preparation protocols for scarred eigenstates and their superpositions that enable their dynamical simulation on quantum computers were proposed. In addition to deterministic methods that made use of the structure of such states and their superpositions as well as their tensor network representations that scale linearly with the system size, it was also found that there were stochastic circuits which can prepare certain target states in constant depth, though with exponentially vanishing post-selection probability. A mix of variational and MPS inspired circuits that were found numerically to construct the target states with great accuracy were employed. The team also provided proof-of-principle state-preparation demonstrations on superconducting quantum hardware.

Quantum many-body scar states are highly excited eigenstates of many-body systems that exhibit atypical entanglement and correlation properties relative to typical eigenstates at the same energy density. Scar states also give rise to infinitely long-lived coherent dynamics when the system is prepared in a special initial state having finite overlap with them. When perturbed, the dynamics of such systems are generally hard to simulate classically, but efficient with quantum computers.

### 4.2.3. Advancements in Hamiltonian Simulation Techniques

The work presented in the paper [[91]] is a significant stride in the realm of Hamiltonian simulations for chemistry applications, addressing the critical issue of gate complexity. Unlike traditional methods that primarily rely on the Hamiltonian itself, leading to approximations that may not fully capture the essential aspects of a problem, SQuISH innovatively uses both the Hamiltonian and an approximate eigenstate. This dual approach enables the creation of a truncated Hamiltonian with reduced complexity, enhancing the efficiency and accuracy of simulations. Integrating this with the ongoing efforts to minimize gate complexity, the Self-consistent Quantum Iteratively Sparsified Hamiltonian (SQuISH) algorithm stands out for its ability to balance computational efficiency with the depth of simulation, making it a pivotal development in quantum computing for chemistry and material science applications [[91]].

### 4.3. Quantum Simulation for Chemistry

In recent years, neural networks have provided an alternate framework for solving electronic structure problems, including applications in directly solving the time-independent Schrödinger equation to find approximate ground states of atoms and molecules (see for example, [92]). A notable example of such a neural network is Fermi Net [93]. In [[94]], Fermi Net was improved with Diffusion Monte Carlo (DMC). Additionally, this work introduced several modifications to the network and optimization approaches that reduce the number of required resources while maintaining or improving the modeling performance.

In related work, [[95]] developed a periodic neural network *Ansatz* for variationally finding the ground-state wavefunction of the homogeneous electron gas (HEG), which is a fundamental model of condensed matter theory. This work extended previous applications of neural networks to molecular systems with methods for handling periodic boundary conditions and includes two notable changes to improve performance: splitting the pairwise streams by alignment and generating backflow coordinates for the orbitals from the network.

A broader perspective on quantum computing for chemical engineering [[96, 97, 11]] introduces concepts for an audience of chemical engineers and discusses three potential first applications for this field: computational chemistry, optimization, and machine learning.

#### 4.3.1. Advances in classical solvers for chemistry Hamiltonians

The electronic structure problem—i.e., the determination of the many-body wavefunction of all electrons in the field of fixed nuclear charges—is the central problem of quantum chemistry and it is most rigorously solved by finding the exact low-lying eigenstates of the electronic many-body Hamiltonian. Of course, because the number of basis states grows exponentially as the number of ways of arranging the electrons in the chosen one-particle orbital basis set, exact solutions are intractable on classical computing hardware for all but the simplest chemical systems. This, and the central importance of this problem, make it an ideal simulation target for quantum computers. Nevertheless, despite the promise/hope of quantum hardware-based solution methods, classical HPC computing methods remain crucial: (i) for benchmarking eventual fault-tolerant quantum hardware on realistic problem instances, and (ii) for preparing initial/starting states on quantum hardware which cannot be too far off from the true solution for the usual quantum algorithm to run efficiently (or at all).

Thus, an ongoing focus of work in the group is to improve state-of-the-art classical approaches to the electronic structure problem, in particular, to develop new HPC-friendly algorithms which leverage the notion that, even for chemistry applications, one does not need to solve the exponentially-scaling eigenproblem to machine precision and, moreover, for initial state preparation, relaxed accuracy requirements may be more acceptable still. In this spirit, a truncated version of the traditional Davidson method for diagonalizing chemistry Hamiltonians has been

developed [[98]] which generates “chemically-accurate” solutions (answers within a chemically-relevant 1.6 milli-Hartrees of the exact solution) for the ground and first few low-lying excited electronic eigenstates at vastly reduced classical computational cost versus a full diagonalization, and that is competitive with other state-of-the-art approximate eigensolvers. Work-in-progress in the group involves developing efficiently parallel distributed-memory versions of the methodology which may be quite effective due to the structure of the newly-developed classical algorithms and their implementation.

#### 4.4. Open Quantum System Simulation

Realistic quantum systems inevitably interact with external degrees of freedom (the environment), which leads to a non-unitary open system evolution. This presents a challenge for the simulation of many relevant physical phenomena on conventional quantum computers that natively implement unitary gates. For this reason, there has been much less research into the emulation of open systems on unitary based quantum computers. In a recent work [[99]], presents an analysis of the resources required for prominent techniques to implement non-unitary operators, some of which incurred a prohibitive classical overhead. Improvements to these algorithms were made, by proposing a fully quantum method without such limitations. Additionally, a quantum algorithm that can decompose any  $d$  dimensional operator into two separate unitaries, which can then be implemented deterministically, was proposed.

#### 4.5. Quantum simulation of High Energy Physics

High energy physics provides an interesting avenue for quantum simulation and possible demonstration of quantum advantage. In this arena, several fundamental aspects of nature are described by quantum field theories, and in particular gauge theories. As a starting point however, it is worthwhile to consider simple quantum field theories, and in [[100]], members of the team investigated how to simulate a simple scalar field on qudit hardware.

While many proposed methods exist for simulating gauge theories describing high energy physics, a computationally promising direction is to approximate continuous gauge groups with discrete subgroups, thereby translating the problem to simulating discrete gauge theories, which are more readily mapped onto digital quantum hardware. Discrete groups are represented by mapping elements of the group to bit strings stored on registers of qubits. It was shown in [101] that 4 primitive group operations are necessary for Trotterized time evolution of pure gauge theories: group inversion, multiplication, phasing, and Fourier transformation. Members of the team derived quantum circuits of the unitary operators for the dihedral groups ( $\mathbb{D}_N$ ) [[102]],  $\mathbb{Z}_2$  lattice gauge theory[[103]], as well as a subgroup of  $SU(2)$ , the binary octahedral group [[104]].

In [[102]], members of the team laid out the framework for constructing quantum gates for discrete groups by enumerating the gates for  $\mathbb{D}_N$ , a simple non-Abelian group. While  $\mathbb{D}_N$  theories map to the group  $U(1) \times \mathbb{Z}_2$  in the large  $N$  limit, which may

by itself be of limited interest, the tools and methods developed extend to future work. In addition to the quantum circuits derived in [[104]], the team found that many of the operators can be decomposed similarly to the classical fast Fourier transform; the operations can be broken down into operations on a series of subgroups. The study also investigated the resource costs for time evolution and compared it to the best reported standard; the team found this representation of an SU(2) gauge theory reduced the required T-gates for simulation by 8 orders of magnitude.

## 5. Tools for Quantum Computing Investigations

Validation and benchmarking of quantum devices are critical for the development of such emerging technologies. Modeling quantum effects in software allows one to verify not only that the behavior is as expected but also to aid in the design of new devices. A prime example of this is a joint work between Google and NASA from 2019 and 2020, respectively, [[105, 106]], which established for the first time a quantum over classical advantage, which used a combination of classical simulation (led by NASA) to verify experimental results. The team continues to develop technologies for quantum simulation in various settings and architectures and also to design methods for benchmarking and characterization.

### 5.1. Classical Methods for Simulating Quantum Circuits

Simulating deep quantum circuits with many qubits is a daunting task. Indeed, while quantum resources only scale polynomially with the number of qubits and the depth of the quantum circuits, classical resources scale exponentially. To overcome this limitation and further raise the bar for a quantum advantage, many classical methods have been developed to tackle different class of quantum circuits. For instance, it is well known that quantum circuits composed of Clifford gates only [107, 108] can be simulated in polynomial time. Indeed, since Clifford gates are nothing else than the stabilizers of the Pauli group, only a polynomial amount of classical resources are required to store the full quantum state. However, since Clifford gates are neither classical universal (that is, not all classical circuits can be described using Clifford gates) nor quantum universal (that is, not all quantum circuits can be described using Clifford gates), their application is limited to error-correction algorithms. Nevertheless, if the quantum circuit is almost composed of Clifford gates, the classical computational cost is exponential in the number of non-Clifford gates only [108, 109].

In [[110]], a technique called Clifford “expansion” was used to classically simulate large out-of-time-order correlators with a few non-Clifford gates, and compare the numerical results with the experimental ones. While the Clifford expansion technique allow the simulation of deep circuits with many qubits with a few non-Clifford gates, the method becomes quickly impractical. While alternative methods that use approximate magic states [109] or noisy gates [111] may help to mitigate the computational cost, Clifford expansion is limited to those quantum circuits with a limited number of non-Cliffords.

On the other hand, methods based on tensor network contraction [112] have been proven to be of practical use to benchmark NISQ devices [113, 84, 114, 115, 116, 117]. The basic idea behind the tensor network contraction is to represent an arbitrary quantum circuits as a network where each gate is a different tensor. Therefore, the calculation of amplitudes reduces to the contraction of the tensor network [112]. While the Clifford expansion method is limited by the number of non-Cliffords, the tensor network contraction is limited by a quantity called “tree-width”, which is closely related to how close the network is to a tree. That is, if the tensor network representing the quantum circuit is tree-like, its tree-width will be close to zero. On the contrary, if the quantum circuit has many entangling gates and its tensor network representation is close to a fully-connected graph, its tree-width will be of the order of the number of qubits. As demonstrated in [112], it is possible to deterministically simulate a quantum circuit in a time that grows exponentially with the tree-width. However, finding the tree-decomposition of a graph that minimize the tree-width is an NP-Complete problem [118]. Moreover, the tree-width does not take into account the fact that, during the contraction, some tensors may exceed the available memory. To mitigate this problem, techniques like slicing have been introduced [116, 113, 115, 117], which add an exponential overhead to the contraction but keep the tensor contraction within the maximum amount of memory available.

Another widely used method to simulate quantum circuits is based on the MPS approximation [119, 120]. An MPS is a quantum state of the form

$$|\psi\rangle = \sum_{\{s\}} \text{Tr} [A_1^{(s_1)} A_2^{(s_2)} \cdots A_N^{(s_N)}] |s_1 s_2 \dots s_N\rangle, \quad (2)$$

where  $A_i^{(s_i)}$  are square matrices of order  $\chi$  (called local dimension), and the indices  $s_i$  go over states in the computational basis. The order  $\chi$  represents how much the qubits are entangled: more precisely, a system of non-entangled qubits can be represented by an MPS with  $\chi = 1$ . On the other hand, a system of highly-entangled qubits will require  $\chi$  being exponentially large with the number of entangled qubits. Since the simulation cost is proportional to  $\chi$ , approximate MPS simulations can be performed by capping the largest allowed  $\chi$  [[113]].

Large-scale quantum computation requires encoding ‘logical qubits’ using an error correction code. Due to a favorable topology, the surface code is a promising candidate for superconducting qubits. One challenge for superconducting qubits, however, is properly isolating the lowest two energy levels of an anharmonic oscillator, with leakage referring to the presence of a population outside of this subspace (which occurs via several physical mechanisms). The simulation of small error correction codes is also critical to aid in processor design, however, each additional level one wishes to simulate incurs an exponential overhead (e.g., for a single leaked state, a fully quantum simulation of  $n$  qubits scales as  $O(3^n)$  instead of  $O(2^n)$ ). In [[121]] it was shown that in error correction codes under realistic situations, an approximation can be employed to faithfully simulate a system with leakage (of any number of leaked states), but still

retaining the  $O(2^n)$  scaling. Given access to HPC resources, this framework enabled a quantum simulation of a distance 5 surface code with leakage. Moreover, this approximation allows one to justify a stronger approximation (Pauli twirling), which can then allow leakage to be simulated *efficiently* (polynomially in  $n$ ) via the stabilizer formalism.

## 5.2. Photonics

A heralded source of highly pure and indistinguishable photons is a prerequisite for many quantum information processing tasks in a photonic (e.g., linear optical) setting. Indeed, the degree of distinguishability is related to the classical hardness of Boson sampling [122], and the ability to generate entangled photonic states. One promising approach to this end is to use photon distillation, first observed in [123] (by Sparrow and Birchall), and improved upon by [124], where linear optics and post-selection is used to filter out erroneous photons in a heralded manner (in contrast to using spectral filters), in principle, to arbitrarily high purity/indistinguishability. This was recently improved in Refs. [125], [126], in which more general families of distillation protocols were given, leading to greater efficiency (i.e., linearly scaling resource requirements).

Another key task for the development of quantum computing primitives is simulation, which is notoriously challenging due to the largeness of the Hilbert space; for  $n$  photons in  $m$  modes, it scales as  $d_{n,m} := \binom{n+m-1}{n}$ . Indeed, the ‘naive’ approach to simulation at the second quantization scales linearly in the dimension. However in many situations much more efficient representations can be found (e.g., in certain cases a polynomially scaling simulation can be performed, similar in spirit to the stabilizer representation [127]). Ref. [128] showed that by realizing Fock states as an approximate superposition of coherent states, a scaling equivalent to qubit systems,  $O(2^n)$ , can be achieved in the linear optical setting, for computing transition amplitudes. Moreover, such an approach is in many cases trivially parallelizable, since each coherent state in the representation can be updated independently of the others.

Unitary designs provide a systematic way to approximate Haar-random unitaries [129] and have a multitude of applications in quantum computing such as randomized benchmarking [130], classical shadow tomography [131], etc. In a recent work by the NASA team, the structure of unitary designs in photonic architectures was investigated [132]. Linear optical unitaries can *only* form a 1-design and are limited in their capabilities for various applications that require  $t \geq 2$ -designs. It was proven that adding a single ‘SNAP’ gate to the group of linear optical unitaries transforms them into a universal gate set for quantum computing and hence allows for the realization of *arbitrary*  $t$ -designs in these systems. However, the structure of linear optics is such that the SNAP gate is a *nonlinear* interaction, which is necessarily probabilistic. Understanding the cost of these nonlinear operations is an interesting research direction for the future.

## 5.3. Quantum Computing Benchmarking

As part of the effort to measure the performance of quantum computing algorithms, it was necessary to provide a set of

benchmarking instances representative of various features of application problems and analysis tools able to synthesize the results in a way that makes them interpretable for users and developers of these computational tools. As highlighted throughout the manuscript, the Hamiltonian formalism provides a well-suited way of incorporating practical applications for quantum computers. Addressing this challenge is also an active area of research that is currently being pursued, and an initial contribution has been the release of the library of Hamiltonians, HamLib [133]. HamLib provides several hundreds of Hamiltonians representing instances of computational problems such as binary variable optimization, including MaxCut, Max-k-Cut, and QMaxCut; discrete variable optimization including the Max-k-SAT and the traveling salesperson problem; condensed matter physics models, including the transverse field Ising model, the Heisenberg model, Fermi-Hubbard model, Bose-Hubbard model; and problems in chemistry such as molecular electronic structure and molecular vibrational structure problems. This work aims to provide researchers with prepared and mapped problem instances, which allow for reproducibility and standardization in research studies and further tests and developments of quantum algorithms.

Among these tests, Ref. [134] showcases a collaborative effort focused on using optimization applications as quantum performance benchmarks. In this work, the canonical instances of MaxCut were considered and addressed via QAOA, implemented in different hardware realizations of qubits, including superconducting and ion qubits and classical simulations, and QA implemented in D-Wave quantum annealers. This work highlights the differences in performance between these algorithms and hardware implementations and provides an example of performance benchmarking of quantum computers by addressing an application.

To facilitate the benchmarking analysis, our team has also been working on implementing analysis tools, drawing inspiration from classical optimization and optimization research. By considering parametrized stochastic optimization solvers as a mathematical framework able to describe QA and QAOA solution methods, among other classical and quantum approaches to optimization, a set of analytical techniques was developed to visually inform about the solver performances and aid in the parameter tuning procedure. The benchmarking method is described in Ref. [135] and its implementation is available in Ref. [136].

## 5.4. Noise Modeling and Characterization

The modeling and characterization of noise in near-term quantum computational devices are critical to aid in the design of such devices, with the aim of achieving fault tolerance. One key aspect to this is using the right metric; whilst full-scale state or process tomography would provide, in principle, all information about a system, it is intractable at modest sizes. It is, therefore, necessary to design metrics that distill the relevant information from the state/process but are still efficiently accessible experimentally. A common metric is the average fidelity; the overlap of the evolved state with the ideal target state, uniformly averaged over all initial states. However, this fidelity is



only a single element in the  $4^n \times 4^n$  process matrix describing the noisy evolution. In [[137]], it was shown that averaging initial states non-uniformly allows for a fidelity metric that is determined by many elements of the transfer matrix. Thus, it can capture additional properties that the standard average fidelity is blind to.

Related to this discussion, in [[138]], it was shown that approximations of global noise channels to local ones allow for an efficient protocol to extract the noise parameters of the process matrix. This work was targeted at repetitive circuits, such as QAOA, where the noise after each round can be estimated, with an experimental demonstration using Rigetti’s quantum processor.

Additionally, noise specifically in QAOA was studied in [[139]], which gave analytical predictions of performance in the presence of local gate noise. This work additionally studied the trade-off between deeper circuits, which ideally do better at optimizing the cost function, but at the expense of more noise acting in the system, which tends to do the opposite.

Another aspect restricting quantum control are frequency fluctuations. In particular, low frequency (such as  $1/f$  or other temporally correlated) noise has been analyzed [[140]] in terms of noise statistics in a repetitive sequence of Ramsey pulses. Moreover, the same Ramsey protocol can be exploited in the nonergodic limit of the noise (shorter than the correlation length) [[141]] to provide supplemental properties of noise, that are otherwise washed out in the ergodic regime. Work led by the Google Quantum AI team [[142]] showed experimentally that non-Gaussianity in low-frequency noise can be observed in superconducting qubits, by performing dynamical decoupling (DD) sequences. Additionally, in [[143]] the QuAIL team worked with Rigetti to use DD to characterize undesired two-qubit couplings as well as the underlying single-qubit decoherence and in suppressing them. A syncopated DD technique is developed that protects against decoherence and selectively targets unwanted two-qubit interactions.

### 5.5. Quantum Error Correction and Mitigation

Quantum error correction (QEC) threshold theorems show that increasing the size (distance) of a QEC code can enable arbitrarily small logical error rates. However, a low enough (code dependent) physical error rate is required in order to guarantee one is in this regime; otherwise the additional qubits can actually cause the logical error rate to *increase*. It was therefore a remarkable result by Google Quantum AI (with assistance in noisy simulations from NASA) in 2023 that showed experimental evidence of a logical error rate reduction by going from a distance 3 to distance 5 surface code [144], a necessary step in the path towards fault tolerance.

One of the remarkable advances in error correction in recent years was the discovery of Floquet codes with dynamical logical qubits. In these codes, the protected subspace changes dynamically, with the logical information updating so as to remain in the protected subspace. The first examples of such codes were honeycomb codes, with the techniques generalizing to trivalent lattices. Recently, members of the QuAIL team developed a Floquet code on a rectangular lattice [[145]].

In [[146]], logical shadow tomography (LST), an error-mitigation technique that is inspired by the QEC theory and more suitable for implementation on near-term quantum computers was developed. LST is a technique to estimate error-mitigated expectation values on noisy quantum computers. The framework is a marriage and generalization of two powerful error mitigating methods: the subspace expansion and virtual distillation. Our technique performs shadow tomography on a logical state to produce a memory-efficient classical reconstruction of the noisy density matrix. Using efficient classical post-processing, one can mitigate errors by projecting a general nonlinear function of the noisy density matrix into the code space. The method was shown to be favorable in the quantum and classical resources overhead. Our technique requires reasonable number of samples,  $O(4^k)$ , to estimate a logical Pauli observable with  $[[N, k]]$  error correction code. Relative to virtual distillation, this technique can compute powers of the density matrix without additional copies of quantum states or quantum memory.

In [[147]] we introduced a new classical shadow tomography protocol based on dual-unitary brick-wall circuits that we termed “dual-unitary shadow tomography” (DUST). We first showed that even at the 2-qubit level, dual-unitaries outperform random Clifford unitaries in that they have a smaller shadow norm (and hence require fewer samples to estimate observables). This is intimately related to the fact that dual-unitary circuits have maximal operator entanglement while random Clifford circuits do not. We generalize this observation to the multiqubit case and by using a transfer matrix method, show that the fast-thermalizing properties of dual-unitary circuits make them better at predicting large operators compared to shallow brick-wall Clifford circuits. In fact, DUST outperforms standard Clifford shadows at any depth  $t$  for predicting extensive observables, and this effect is even more pronounced for small system sizes.

In recent work [[148]], an error mitigation method for certain realistic noise models that can be applied given partial knowledge of noise to achieve lower biases than typically possible with other methods was presented. Subspaces that decay uniformly were obtained and used to perform error mitigated quantum computation. The expectation values of dynamics encoded in such subspaces are unbiased estimators of noise-free expectation values. This theory was applied to a system of qubits and qudits undergoing relaxation with varying decay rates and show that such subspaces can be used to obtain unbiased estimation of the true expectation value up to second order variations in the decay rates. Given varying decay rates, this method was implemented in analog and circuit-model settings, improving on dual-rail qubits, and given partial knowledge of noise, show that it can outperform probabilistic error cancellation.

### 5.6. Pulse Control

Together with colleagues at the University of Michigan and Argonne National Lab, the application of various classical optimization techniques to pulsed or bang-bang controls for quantum problems was studied. This work originally focused on

various classical methods applied to the quantum control setting [[149]] but has been expanded to considering new heuristics for reducing the number of control switches [[150]] and to considering unitary control errors [[151]]. This work all relies on a continuous relaxation of the control problem, followed by various techniques for rounding and fixing the controllers to be binary pulses, akin to what would appear in QAOA. Constraints on the length of the pulses and the switching time are also considered and incorporated into this framework. Combinations of various classical techniques are used in attempts to improve on performance while satisfying physical and noise constraints on the system. These techniques are applied to toy problems, such as minimizing the energy of a transverse field Ising model, compiling quantum circuits, and constructing quantum gates from physical controllers.

In [[152]], the problem of co-designing algorithms on novel Circuit Quantum Electrodynamics (cQED) architectures enabled by 3D superconducting radio-frequency cavities such as those developed at SQMS [[153]] is tackled numerically. This is accomplished by setting up a pulse-engineering approach to compiling the primitive gates of QAOA in qudit systems. The customized use of the `juqbox` quantum optimal control package allowed the synthesis of a single round of the Max-k-Cut QAOA algorithm for arbitrary parameters at high-fidelity, on a realistic noiseless model describing the Hamiltonian of one or two cavity modes. In subsequent work, other SQMS researchers generalized the model including the effect of additional modes [154].

### 5.7. Generating Hard Problem Instances

With the development of more competitive classical [71], quantum [105, 113, 155], or “quantum-inspired” [156, 157] heuristics, it is becoming of utmost importance to identify classes of instances which are well suited to transversally benchmark such devices. While it is well known that random instances have a hardness threshold [158, 159, 160, 161], it is a major challenge to designing planted problems with known solutions which are hard to optimize: indeed, planting a desired solution may affect the energy landscape of the instance, making it more convex and therefore easier to explore. In the past few years, multiple planting methods have been developed, including “quiet” planting techniques [162, 163, 164, 165, 166, 167], “patch” planting [168, 169], and the use of cryptographic protocols [[170]]. Hard instances can also be generated from more practical applications such as planning problems [171].

Ref. [[170]] proposed a novel way of randomly generating hard instances of disordered Ising Hamiltonians with unique planted ground states by casting the public key of the McEliece cryptographic protocol [172] as an interacting system of Ising spins, so that finding the ground state of the resulting Hamiltonian is equivalent to breaking the associated McEliece cryptosystem by only knowing its public key, while the private key allows a party who knows it to recover the planted ground state. Thanks to this equivalence, the security of the McEliece cryptosystem in the typical case is leveraged in order to produce Ising Hamiltonians whose ground state is computationally hard

to find. Secondly, by using mathematical methods originating in the statistical physics of disordered systems (see [173] for an accessible introduction) the authors have studied in simplified settings the effect of random energy scrambling on easy-to-optimize energy landscapes, following the notion that in the McEliece cryptosystem the public and the private keys are connected through random permutations. It was shown that the energy-scrambling can introduce a phenomenon known as “clustering” into the energy landscape of the Hamiltonian: close enough to the ground-state, the states with a given energy density organize themselves in well-separated groups (i.e. “clusters”), indicating a complicated geometry for the energy landscape of the kind that common local-search heuristics find it difficult to explore.

In an effort to provide a unified cross-platform framework for generating random instances with planted solutions, Ref. [[174]] devised `chook`, an open-source and user-friendly library to generate random instances with known solutions and tunable hardness. `chook` is designed to be extensible and updated to include new developed planting methods. At the time of writing, `chook` includes generators such as the Wishart planting [175], equation planting [176], deceptive cluster loop planting [173], as well as planted solutions for higher-order (beyond quadratic) binary optimization problems.

### 5.8. Randomized Benchmarking

Randomized benchmarking (RB) is a powerful method for determining the error rate of experimental quantum gates [130]. As a resource-efficient alternative method to quantum process tomography, RB also provides an estimate for the average fidelity that is independent of state preparation and measurement (SPAM) errors. RB characterizes gates that form a group  $G$  that form a unitary 2-design in a resource-scalable fashion. Character RB can benchmark more general gates using techniques from representation theory but has only been explored for “multiplicity-free” groups. In the work of [[177]], the character RB theory to explicitly treat non-multiplicity-free groups was developed. Our rigorous derivations enabled us to provide conditions under which instantiations of the framework yield practical RB protocols. This generalized approach with applications to three distinct situations of practical interest: benchmarking of gates with subspace-preserving properties, characterization of leakage, and benchmarking of the match-gate group was demonstrated. In all three cases, compared to existing theories, our method requires similar resources but either provides a more accurate estimate of gate fidelity or applies to a more general group of gates.

While the structure of unitary 2-designs for multiqubit systems is well understood by now, little is known about exact 2-designs for *multiqudit* systems, especially for non-prime-power dimensions [178, 179, 180]. In an upcoming work, a class of weighted state 2-designs for qudits in arbitrary dimensions [[181]] that allow us to devise several new protocols for RB, shadow tomography, etc., in various qudit architectures.

### 5.9. Master Equation Emulation

Prior work by Marshall et al. in [[182]] led to an experimental implementation from the Levenson-Falk group at USC, showing a high degree of control of a non-Markovian master equation using superconducting qubits [[183]]. It was further demonstrated, as predicted in [[182]], that introducing non-Markovian noise to a Markovian bath can actually, somewhat counter intuitively, *increase* the coherence time of a qubit.

## 6. Mechanisms for Quantum Computation

This section describes a broad set of topics of a more advanced nature, often of a more fundamental nature all generally related to quantum information and its ties to quantum computation. Most of this work is more theoretical in nature, such as work on the spread of information through a quantum system or on how basis dependent properties interact during quantum evolution. There are also more practical focuses such as how to design logic gates and to translate quantum effects into equivalent classical enhancements and architectures for large scale quantum computing. We begin with a discussion of work related to fundamentals of quantum theory and the use of a quantum computer in testing the boundaries of quantum theory.

### 6.1. Wigner Friend Scenario

Bell’s theorem enable tests of quantum theories that later underpinned quantum information processing. As mentioned in Sec. 2.2, it gives rise to entanglement witnesses which are one way of investigating quantum entanglement in quantum states, including those arising at each stage of an algorithm or protocol. Recently, researchers realized that the Wigner friend scenario, an equally intriguing but less known thought experiment than Schrödinger’s cat, can be extended to give rise to new inequalities, local friendliness inequalities, that have weaker assumptions than those in Bell’s theorem. These inequalities provide new experimental horizons, extending from an already realized proof-of-principle experiment in which a single photon plays the role of Wigner’s friend to a full experiment involving a quantum computer running a human-level artificial intelligence with a requirement of space-like separation between the parties of the distance between the moon and earth or more. In [[184]], Rieffel together with two Australian colleagues estimated the resources for this full experiment and helped refine the statement of a new “local friendliness” theorem. While both a sufficiently large and robust quantum computer and a human-level AI have not been realized today, they appear technically feasible within the next decades. The full experiment provides a challenging target, with related nearer-term feasible experiment targets that test the boundaries of quantum theory being explored.

### 6.2. Information Scrambling

Information scrambling is a fundamental mechanism by which closed, interacting quantum many-body systems spread information throughout their nonlocal degrees of freedom

[185, 186]. Scrambling phenomena often underlies methods to demonstrate and benchmark quantum computational capabilities. While typically characterized via the decay of “out-of-time-ordered correlators” (OTOCs) [187], this delocalization manifests itself via two distinct processes: (i) “operator spreading,” which corresponds to the increasing support of local operators under Heisenberg evolution [188] and (ii) “operator entanglement generation” which represents the increasing complexity of Heisenberg-evolved operators, quantified via the inability to express them as a simple tensor products of local operators [189].

The dynamics of operator spreading and operator entanglement growth have turned out to be crucial in understanding minimal models of quantum chaos such as “dual-unitary circuits” [115] and “Sachdev-Ye-Kitaev” (SYK) type models [187], as well as simple models for the black hole information paradox such as the Hayden-Preskill protocol [190]. Beyond these somewhat exotic applications, these two mechanisms are also intimately related to classical simulability of quantum dynamics [191]. For example, in local Hamiltonian systems on a lattice, operator spreading is controlled via Lieb-Robinson (LR) bounds [192], which provides an estimate for the maximal locality one needs to consider in order to approximate Heisenberg-evolved operators well. In Ref. [[110]], LR bounds were used to introduce a “lightcone filter” which allowed them to improve experimental accuracy in measuring OTOCs. In contrast, operator entanglement is a nonlocal observable that does not directly follow such a (exponential) bound. However, in the same spirit as above, the growth of operator entanglement is directly related to the simulability of Heisenberg operators via a matrix product operator (MPO) representation [191]. This simulation cost shows up in tensor network contraction approaches to classically simulating quantum dynamics [[193]].

In a work led by the Google Quantum AI team, both operator spreading and operator entanglement growth were observed in a 53-qubit superconducting quantum processor [[110]]. This demonstrated a high degree of control, verified by classical simulations from the NASA QuAIL team. They considered two different classes of random unitaries (by either using SWAP gates or  $\sqrt{\text{SWAP}}$  gates) to experimentally realize different rates of operator spreading and operator entanglement growth. Furthermore, they introduced models to simulate the growth of OTOCs in large-scale systems. A key observation here was that the classical resources required to simulate operator growth are much tamer compared to those required for simulating operator entanglement, consistent with previous numerical and analytical studies.

These experiments have motivated further study on the connection between OTOCs and operator entanglement, which was rigorously established in Ref. [[194]]. A series of works generalized this connection to open quantum systems [[195]] and systems at finite temperature [[196]], and placed it in an operator algebraic framework [[197]]. Building on these theoretical and experimental advances, the QuAIL team along with researchers at USC, further generalized these results to include non-Hermitian Hamiltonians, which typically violate LR bounds, and represent effective models for



measurement-induced phase transitions [[198]].

### 6.3. Stoquasticity

Stoquasticity and the sign problem are related to the difficulty of classically simulating quantum phenomena. A problem or Hamiltonian with no sign problem is simulatable using classical Monte Carlo techniques, but when a sign problem occurs, the pseudo-probabilities in the quantum Monte Carlo technique start oscillating rapidly between positive and negative, leading to an exponential slow down in the integration/averaging. The sign problem occurs in most quantum Hamiltonians, but stoquastic Hamiltonians are a notable class where no sign problem occurs. Stoquasticity is a basis-dependent property of the Hamiltonian and is characterized by the Hamiltonian being represented by a matrix whose off-diagonal entries are all non-positive.

When a single Hamiltonian is being simulated, a stoquastic basis always exists (e.g., the eigenbasis of the Hamiltonian), but finding a stoquastic (or more generally a sign problem free) basis, also known as “curing” the sign problem is an NP-hard task [199]. However, some quantum algorithms, such as quantum annealing and QAOA are characterized by the interplay and evolution under two distinct and non-commuting Hamiltonians. In Ref. [[200]], the question of whether a single basis exists under which multiple Hamiltonians are stoquastic was explored. If such a basis, coined the “simultaneous stoquastic” basis exists, then it would be possible to simulate a quantum annealing or QAOA task using Monte Carlo techniques with no sign problem. Our results, based off the theory of unitary similarity, show that for any given set of Hamiltonians of size  $d \geq 2$  and dimension  $N > 2$ , the probability that a simultaneously stoquastic basis exists is vanishingly small. Furthermore, a simple no go result for testing this based off the commutator of the Hamiltonians was provided. This could potentially have further applications in the computational complexity of simulating such Hamiltonian processes where no such basis exists.

### 6.4. Relaxation Mechanisms in Quantum and Classical Transverse-Field Annealing

Prior work in Ref. [[201]] showed that strategically introducing a pause in a quantum annealing schedule can increase the performance by several orders of magnitude. In that work, it was posited that thermal relaxations were occurring during the pause after excitations from driving through the minimum gap. A follow-up study in [[202]] showed that the same behavior could be found in a purely classical analog of quantum annealing (SVMC), thus providing more insights into the origin of this effect (i.e., it is not a uniquely quantum phenomenon). This work also gave a general condition on pausing for it to yield an improvement agnostic to the particular system.

### 6.5. Acceleration of Quantum Tunneling in Spin Systems

Expanding on the original work on Random-Frequency Quantum Annealing (RFQA) [203], the QuAIL team and their collaborators have explored ways of mitigating the decay of performance in currently-available quantum annealers caused

by minor embedding. The freezing of the physical qubit chains that represent logical qubits is suppressed by adding independently oscillating local fields with a distribution of frequencies to the usual transverse-field annealing setup. This off-equilibrium process allows a chain to escape its frozen state via a quantum tunneling event induced by the proliferation of weak resonances introduced by the oscillatory fields, thus reducing the effect of frozen chains. This was studied in [[204]] using the 1D transverse-field Ising chain as a model for a chain of spins used in the embedding of logical problems. A follow-up work [[205]] designed to shed light on the physics behind this process showed how the same results can be engineered in a driven Floquet system, where additional fields oscillate with the same frequency. Moreover, it uncovered the existence of some novel off-equilibrium phenomenology of these physical systems. Using the Transverse-Field Ising models (TFIM) in 1D and 2D as testing grounds, the authors showed three off-equilibrium regimes as a function of the drive strength: for weak drives, the system exhibits exponentially decaying tunneling rates but robust magnetic order, typical of the (undriven) dynamics in the ordered phase of models exhibiting spontaneous symmetry breaking such as the TFIM. In the crossover regime at intermediate drive strength, one observes polynomial decay of tunneling alongside vanishing magnetic order, and at very strong drive strengths both the Rabi frequency and time-averaged magnetic order are approximately constant with increasing system size.

### 6.6. Engineering Topology

Exotic topological states of matter such as the non-Abelian anyons exhibit braiding statistics foundational to topological fault-tolerant quantum computation and can emerge in systems displaying the fractional quantum Hall effect (FQHE). Topologically non-trivial bands when made flat (degenerate) can lead to FQHE states. In recent work by Suri et al. [[206]], a method to create bands with higher order topology which is not possible to obtain otherwise for example by a simple magnetic field was proposed. Such higher order topological bands when made flat can lead to potentially more exotic non-Abelian anyons useful for quantum computation.

### 6.7. Compilation approaches

Compiling quantum circuits to realistic hardware architectures generally involves two steps, circuit routing and gate synthesis. In [[207]], earlier work, [[208]], was extended, which recognized that routing could be viewed as temporal planning task, and thus state-of-the-art temporal planners applied, to route circuits that implement QAOA for Graph Coloring problems. On the gate synthesis side, Ref. [[209]], studied the case of single qubit quantum compilation and state preparation by modeling the task as a Markov Decision Process, and solving it using dynamical programming methods. This work demonstrates that classical reinforcement learning techniques can produce shorter gate sequences than have been identified before, and can also generate gate sequences that work well in the presence of noise. These results represent an



important advance in methods for quantum compilation that work in the presence of the physical constraints of NISQ-era hardware.

### 6.8. Quantum Computing Architectures

Quantum algorithms are usually framed in terms of qubits in the gate based circuit model. Physical realizations of qubits take the form of several different implementations, such as superconducting hardware, trapped ions, neutral atoms etc. Recently, advances on the experimental side have opened the avenue to control the quantum state of superconducting cavity modes coupled to a qubit-like transmon. In the Fock basis, this in turn enables the control and manipulation of qudits, larger dimensional analogues of qubits. In [[153]], it was describe how such a platform may be realized through 3D superconducting radio frequency (SRF) cavities, which typically have large Q factors. Several algorithms of interest in high energy physics (HEP) involve bosonic degrees of freedom, which are often more easily mapped onto qudits rather than qubits.

A potentially promising approach for assembling future large scale-quantum computers is discussed in [[210]]. Here a modular solid state architecture with deterministic inter-module coupling between four physically separate integrated chips is described, and the quality of the inter-module entanglement is confirmed by a demonstration of Bell-inequality violation for disjoint pairs of entangled qubits across the four separate silicon dies.

## 7. Future Outlook

Many industry, government, and academic groups have ambitious programs to increase the size and improve the robustness of quantum processors. For the next several years, however, a gap will remain between the capabilities of the hardware and the capabilities needed to run quantum and quantum-classical algorithms at application-scale. While performing spectacular one-off experiments are tempting, the field will advance by prioritizing experiments and theory that inform efforts to reach application scale in a robust, repeatable way. The challenges that remain are significant enough to require broad collaborations and open sharing of information.

As the hardware matures, we will have increasing opportunities for exploration of algorithms, significantly beyond what has been possible up to now. For many of the most computationally intensive state-of-the-art practical algorithms, performance must be evaluated empirically – theoretical analysis is just too hard. For this reason, there are machine learning, SAT, and planning and scheduling competitions with sets of challenge problems on which algorithms can compete empirically, and within industry and government, heuristic algorithms, without proofs of their performance, are frequently used to attack practical problems. Currently, evaluation of the performance of quantum algorithms is restricted to small cases, given the size and non-robustness of current quantum processors and that generically there is exponential overhead in simulating quantum algorithms on classical computers, quickly exceeding the

capabilities of even the world’s largest supercomputers. While the size and robustness of quantum processors are expected to improve steadily in the coming years, they will remain resource limited in the next decades. Algorithm-hardware co-design will be critical to reap the most benefit from hardware advances.

Co-design is also needed between hardware and error mitigation, error correction, native gate choices, and compilation of quantum algorithms to native gates. Next generation quantum systems will provide information that will enable richer error models. These models will inform error correction which can in turn inform hardware design priorities and architecture choices. As quantum computers scale up, they will necessarily involve multiple QPUs, first resembling clusters and then supercomputers, but with quantum as well as classical processors, where the classical processors will not only be used for control of quantum processors, but also for decoding as part of error correction, and often as co-processors in hybrid quantum-classical algorithms. We expect research into future HPC systems and quantum systems to mutually inform each other. Further research directions include understanding the space of architectural designs and co-designing distributed algorithms, error correction layers, and classical and quantum communication and synchronization with architectural choices. It is an open question how heterogeneous these quantum computers will be. It is possible that they will include multiple types of quantum processors, superconducting, ion-trap, photonic, neutral atom, etc., and will likely include a variety of specialized quantum components such as quantum memory or magic state factories or special-purpose quantum processors for, say, quantum Fourier transforms. Specialized quantum processors that have less stringent requirements than universal quantum computers could potentially be built more easily and put to use earlier than universal processors. Resource estimation for specific problems of application interest, including resource estimates for newly considered problems to improved estimates for already identified problems, will help guide hardware, error correction, and architectural priorities.

Quantum computing is one of the most enticing computational paradigms, with the potential to revolutionize diverse areas of future-generation computational systems. The NASA QuAIL team looks forward to continuing to advance the understanding of the potential of these systems in collaborations with others in this diverse field. We hope this glimpse of the variety of the current research and the many open areas will further entice both young people at the beginning of their careers and professionals in adjacent fields to join this fascinating effort.

## 8. Acknowledgements

The authors are grateful for our many collaborators who have made much of this work possible. We are grateful for support from NASA Ames Research Center, NASA Advanced Exploration systems (AES) program, NASA Earth Science Technology Office (ESTO), NASA Transformative Aeronautic Concepts Program (TACP), and NASA Space Communications and Navigation (SCaN). We are grateful for the support of the

U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Superconducting Quantum Materials and Systems Center (SQMS) under the contract No. DE-AC02-07CH11359 and Co-Design Center for Quantum Advantage (C2QA) under Contract No. DE-SC0012704, from DARPA under IAA 8839, Annex 114, 128, 129, 130, SAA2-403688, and SAA2-403746, from Google LLC under SSA-403706, and from NTT Research Inc. under SAA2-403506. A.A.A., M.S.A, D.E.B.N., S.B., E.G., Z.G.I, S.H., P.A.L, F.M., J.M., N.S., D.V., and Z.W. are supported by NASA Academic Mission Services (NAMS), contract number NNA16BD14C. N.A., H.M.B, L.T.B., S.C., S.M., and G.M. are KBR employees working under the Prime Contract No. 80ARC020D0010 with NASA.

## 9. About the NASA QuAIL Team

The mandate of the NASA’s Quantum Artificial Intelligence Laboratory (QuAIL), located at NASA Ames in the heart of Silicon Valley, is to assess and advance the potential impact of quantum computers on computational problems that will be faced by NASA missions in decades to come. QuAIL team members, physicists, computer scientists, mathematicians, and engineers, come from a wide variety of backgrounds with complementary expertise that supports the interdisciplinary nature of quantum computing research. The team has strong collaborations with application domain experts in and outside NASA. QuAIL is a theory and numerical group that has forged close collaborations with groups implementing quantum hardware.

## References

[1] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* 2 (2018) 79. URL: <https://doi.org/10.22331/q-2018-08-06-79>. doi:10.22331/q-2018-08-06-79.

[2] E. G. Rieffel, S. Hadfield, T. Hogg, S. Mandrà, J. Marshall, G. Mossi, B. O’Gorman, E. Plamadeala, N. M. Tubman, D. Venturelli, W. Vinci, Z. Wang, M. Wilson, F. Wudarski, R. Biswas, From Ansätze to Z-Gates: A NASA View of Quantum Computing, IOS Press, 2019. URL: <http://dx.doi.org/10.3233/APC190010>. doi:10.3233/apc190010.

[3] R. Biswas, Z. Jiang, K. Kechezhi, S. Knysch, S. Mandrà, B. O’Gorman, A. Perdomo-Ortiz, A. Petukhov, J. Realpe-Gómez, E. Rieffel, D. Venturelli, F. Vasko, Z. Wang, A NASA perspective on quantum computing: Opportunities and challenges, *Parallel Computing* 64 (2017) 81–98. URL: <http://dx.doi.org/10.1016/j.parco.2016.11.002>. doi:10.1016/j.parco.2016.11.002.

[4] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, D. Preda, A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem, *Science* 292 (2001) 472–475. URL: <https://www.science.org/doi/abs/10.1126/science.1057726>. doi:10.1126/science.1057726.

[5] E. Farhi, J. Goldstone, S. Gutmann, A Quantum Approximate Optimization Algorithm, 2014. [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).

[6] E. Farhi, J. Goldstone, S. Gutmann, A Quantum Approximate Optimization Algorithm Applied to a Bounded Occurrence Constraint Problem, 2015. [arXiv:1412.6062](https://arxiv.org/abs/1412.6062).

[7] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, P. J. Coles, Variational quantum algorithms, *Nature Reviews Physics* 3 (2021) 625–644. URL: <https://doi.org/10.1038/s42254-021-00348-9>. doi:10.1038/s42254-021-00348-9.

[8] A. Montanaro, Quantum speedup of branch-and-bound algorithms, *Phys. Rev. Res.* 2 (2020) 013056. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.2.013056>. doi:10.1103/PhysRevResearch.2.013056.

[9] C.-M. Alexandru, E. Bridgett-Tomkinson, N. Linden, J. MacManus, A. Montanaro, H. Morris, Quantum speedups of some general-purpose numerical optimisation algorithms, *Quantum Science and Technology* 5 (2020) 045014. URL: <https://dx.doi.org/10.1088/2058-9565/abb003>. doi:10.1088/2058-9565/abb003.

[10] A. Abbas, A. Ambainis, B. Augustino, A. Bäertschi, H. Buhman, C. Cofrin, G. Cortiana, V. Dunjko, D. J. Egger, B. G. Elmegreen, et al., Quantum Optimization: Potential, Challenges, and the Path Forward, 2023. [arXiv:2312.02279](https://arxiv.org/abs/2312.02279).

[11] D. E. Bernal Neira, C. D. Laird, L. R. Lueg, S. M. Harwood, D. Trenev, D. Venturelli, Utilizing modern computer architectures to solve mathematical optimization problems: A survey, *Computers & Chemical Engineering* (2024) 108627. doi:<https://doi.org/10.1016/j.compchemeng.2024.108627>.

[12] S. Hadfield, Z. Wang, B. O’Gorman, E. Rieffel, D. Venturelli, R. Biswas, From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz, *Algorithms* 12 (2019) 34. URL: <http://dx.doi.org/10.3390/a12020034>. doi:10.3390/a12020034.

[13] J. Sud, S. Hadfield, E. Rieffel, N. Tubman, T. Hogg, Parameter-setting heuristic for the quantum alternating operator ansatz, *Physical Review Research* 6 (2024) 023171.

[14] Z. Wang, S. Hadfield, Z. Jiang, E. G. Rieffel, Quantum approximate optimization algorithm for MaxCut: A fermionic view, *Physical Review A* 97 (2018) 022304. URL: <https://link.aps.org/doi/10.1103/PhysRevA.97.022304>. doi:10.1103/PhysRevA.97.022304.

[15] K. Marwaha, S. Hadfield, Bounds on approximating Max  $k$ XOR with quantum and classical local algorithms, *Quantum* 6 (2022) 757. URL: <https://doi.org/10.22331/q-2022-07-07-757>. doi:10.22331/q-2022-07-07-757.

[16] J. Hirvonen, J. Rybicki, S. Schmid, J. Suomela, Large cuts with local algorithms on triangle-free graphs, 2014. [arXiv:1402.2543](https://arxiv.org/abs/1402.2543).

[17] S. Sen, Optimization on sparse random hypergraphs and spin glasses, *Random Structures & Algorithms* 53 (2018) 504–536. doi:<https://doi.org/10.1002/rsa.20774>.

[18] S. Bravyi, A. Kliesch, R. Koenig, E. Tang, Obstacles to Variational Quantum Optimization from Symmetry Protection, *Phys. Rev. Lett.* 125 (2020) 260505. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.260505>. doi:10.1103/PhysRevLett.125.260505.

[19] H. Leipold, F. M. Spedalieri, E. Rieffel, Tailored Quantum Alternating Operator Ansatzes for Circuit Fault Diagnostics, *Algorithms* 15 (2022). URL: <https://www.mdpi.com/1999-4893/15/10/356>. doi:10.3390/a15100356.

[20] S. Hadfield, T. Hogg, E. G. Rieffel, Analytical framework for quantum alternating operator ansätze, *Quantum Science and Technology* 8 (2022) 015017. URL: <https://dx.doi.org/10.1088/2058-9565/aca3ce>. doi:10.1088/2058-9565/aca3ce.

[21] V. Kremenetski, A. Apte, T. Hogg, S. Hadfield, N. M. Tubman, Quantum Alternating Operator Ansatz (QAOA) beyond low depth with gradually changing unitaries, 2023. [arXiv:2305.04455](https://arxiv.org/abs/2305.04455).

[22] V. Kremenetski, T. Hogg, S. Hadfield, S. J. Cotton, N. M. Tubman, Quantum Alternating Operator Ansatz (QAOA) Phase Diagrams and Applications for Quantum Chemistry, 2021. [arXiv:2108.13056](https://arxiv.org/abs/2108.13056).

[23] R. LaRose, E. Rieffel, D. Venturelli, Mixer-phaser Ansatzes for quantum optimization with hard constraints, *Quantum Machine Intelligence* 4 (2022) 17. URL: <https://doi.org/10.1007/s42484-022-00069-x>. doi:10.1007/s42484-022-00069-x.

[24] F. B. Maciejewski, S. Hadfield, B. Hall, M. Hodson, M. Dupont, B. Everett, J. Sud, M. S. Alam, Z. Wang, S. Jeffrey, B. Sundar, P. A. Lott, S. Grabbe, E. G. Rieffel, M. J. Reagor, D. Venturelli, Design and execution of quantum circuits using tens of superconducting qubits and thousands of gates for dense Ising optimization problems, 2023. [arXiv:2308.12423](https://arxiv.org/abs/2308.12423).

[25] F. B. Maciejewski, J. Biamonte, S. Hadfield, D. Venturelli, Improving Quantum Approximate Optimization by Noise-Directed Adaptive Remapping, *arXiv preprint arXiv:2404.01412* (2024). doi:10.48550/arXiv.2404.01412.

- [26] Z. Wang, N. C. Rubin, J. M. Dominy, E. G. Rieffel, XY mixers: Analytical and numerical results for the quantum alternating operator ansatz, *Phys. Rev. A* 101 (2020) 012320. URL: <https://link.aps.org/doi/10.1103/PhysRevA.101.012320>. doi:10.1103/PhysRevA.101.012320.
- [27] M. Streif, M. Leib, F. Wudarski, E. Rieffel, Z. Wang, Quantum algorithms with local particle-number conservation: Noise effects and error correction, *Phys. Rev. A* 103 (2021) 042412. URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.042412>. doi:10.1103/PhysRevA.103.042412.
- [28] M. S. Alam, F. A. Wudarski, M. J. Reagor, J. Sud, S. Grabbe, Z. Wang, M. Hodson, P. A. Lott, E. G. Rieffel, D. Venturelli, Practical Verification of Quantum Properties in Quantum-Approximate-Optimization Runs, *Phys. Rev. Appl.* 17 (2022) 024026. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.17.024026>. doi:10.1103/PhysRevApplied.17.024026.
- [29] L. T. Brady, S. Hadfield, Iterative Quantum Algorithms for Maximum Independent Set: A Tale of Low-Depth Quantum Algorithms, 2023. [arXiv:2309.13110](https://arxiv.org/abs/2309.13110).
- [30] M. Dupont, B. Evert, M. J. Hodson, B. Sundar, S. Jeffrey, Y. Yamaguchi, D. Feng, F. B. Maciejewski, S. Hadfield, M. S. Alam, et al., Quantum-enhanced greedy combinatorial optimization solver, *Science Advances* 9 (2023) eadi0487. doi:10.1126/sciadv.adi0487.
- [31] M. Dupont, B. Sundar, B. Evert, D. E. B. Neira, Z. Peng, S. Jeffrey, M. J. Hodson, Quantum Optimization for the Maximum Cut Problem on a Superconducting Quantum Computer, 2024. [arXiv:2404.17579](https://arxiv.org/abs/2404.17579).
- [32] T. Kadowaki, H. Nishimori, Quantum annealing in the transverse Ising model, *Physical Review E* 58 (1998) 5355–5363. URL: <https://doi.org/10.1103/PhysRevE.58.5355>. doi:10.1103/PhysRevE.58.5355.
- [33] E. Farhi, J. Goldstone, S. Gutmann, M. Sipser, Quantum Computation by Adiabatic Evolution, [arxiv \(2000\)](https://arxiv.org/abs/quant-ph/0001106). URL: <https://arxiv.org/abs/quant-ph/0001106>. [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
- [34] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, O. Regev, Adiabatic Quantum Computation Is Equivalent to Standard Quantum Computation, *SIAM Review* 50 (2008) 755–787. URL: <https://doi.org/10.1137/080734479>. doi:10.1137/080734479.
- [35] Z. Gonzalez Izquierdo, S. Grabbe, S. Hadfield, J. Marshall, Z. Wang, E. Rieffel, Ferromagnetically Shifting the Power of Pausing, *Phys. Rev. Appl.* 15 (2021) 044013. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.15.044013>. doi:10.1103/PhysRevApplied.15.044013.
- [36] Z. Gonzalez Izquierdo, S. Grabbe, H. Idris, Z. Wang, J. Marshall, E. Rieffel, Advantage of Pausing: Parameter Setting for Quantum Annealers, *Phys. Rev. Appl.* 18 (2022) 054056. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.18.054056>. doi:10.1103/PhysRevApplied.18.054056.
- [37] B. Pokharel, Z. G. Izquierdo, P. A. Lott, E. Strbac, K. Osiewalski, E. Papatheasiou, A. Kondratyev, D. Venturelli, E. Rieffel, Inter-generational comparison of quantum annealers in solving hard scheduling problems, *Quantum Information Processing* 22 (2023) 364. doi:10.1007/s11128-023-04077-z.
- [38] J. Marshall, A. Di Gioacchino, E. G. Rieffel, Perils of embedding for sampling problems, *Phys. Rev. Res.* 2 (2020) 023020. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.2.023020>. doi:10.1103/PhysRevResearch.2.023020.
- [39] J. Marshall, G. Mossi, E. G. Rieffel, Perils of embedding for quantum sampling, *Phys. Rev. A* 105 (2022) 022615. URL: <https://link.aps.org/doi/10.1103/PhysRevA.105.022615>. doi:10.1103/PhysRevA.105.022615.
- [40] Z. Gonzalez Izquierdo, R. Zhou, K. Markström, I. Hen, Discriminating nonisomorphic graphs with an experimental quantum annealer, *Phys. Rev. A* 102 (2020) 032622. URL: <https://link.aps.org/doi/10.1103/PhysRevA.102.032622>. doi:10.1103/PhysRevA.102.032622.
- [41] S. Knysh, E. Plamadeala, D. Venturelli, Quantum annealing speedup of embedded problems via suppression of Griffiths singularities, *Phys. Rev. B* 102 (2020) 220407. URL: <https://link.aps.org/doi/10.1103/PhysRevB.102.220407>. doi:10.1103/PhysRevB.102.220407.
- [42] D. E. Bernal, K. E. Booth, R. Dridi, H. Alghassi, S. Tayur, D. Venturelli, Integer programming techniques for minor-embedding in quantum annealers, in: *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 17th International Conference, CPAIOR 2020, Vienna, Austria, September 21–24, 2020, Proceedings 17*, Springer, 2020, pp. 112–129.
- [43] C. M. Unsal, L. T. Brady, Quantum Adversarial Learning in Emulation of Monte-Carlo Methods for Max-cut Approximation: QAOA is not optimal, 2022. [arXiv:2211.13767](https://arxiv.org/abs/2211.13767).
- [44] L. P. García-Pintos, L. T. Brady, J. Bringewatt, Y.-K. Liu, Lower Bounds on Quantum Annealing Times, *Phys. Rev. Lett.* 130 (2023) 140601. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.140601>. doi:10.1103/PhysRevLett.130.140601.
- [45] T. Stollenwerk, S. Hadfield, Z. Wang, Toward Quantum Gate-Model Heuristics for Real-World Planning Problems, *IEEE Transactions on Quantum Engineering 1* (2020) 1–16. doi:10.1109/TQE.2020.3030609.
- [46] T. Stollenwerk, B. O’Gorman, D. Venturelli, S. Mandrà, O. Rodionova, H. Ng, B. Sridhar, E. G. Rieffel, R. Biswas, Quantum Annealing Applied to De-Conflicting Optimal Trajectories for Air Traffic Management, *IEEE Transactions on Intelligent Transportation Systems* 21 (2020) 285–297. doi:10.1109/TITS.2019.2891235.
- [47] K. E. Booth, B. O’Gorman, J. Marshall, S. Hadfield, E. Rieffel, Quantum-accelerated global constraint filtering, in: *Principles and Practice of Constraint Programming: 26th International Conference, CP 2020, Louvain-la-Neuve, Belgium, September 7–11, 2020, Proceedings 26*, Springer, 2020, pp. 72–89. doi:10.1007/978-3-030-58475-7\_5.
- [48] K. E. C. Booth, B. O’Gorman, J. Marshall, S. Hadfield, E. Rieffel, Quantum-accelerated constraint programming, *Quantum* 5 (2021) 550. URL: <https://doi.org/10.22331/q-2021-09-28-550>. doi:10.22331/q-2021-09-28-550.
- [49] R. Shaydulin, S. Hadfield, T. Hogg, I. Safro, Classical symmetries and the Quantum Approximate Optimization Algorithm, *Quantum Information Processing* 20 (2021). URL: <http://dx.doi.org/10.1007/s11128-021-03298-4>. doi:10.1007/s11128-021-03298-4.
- [50] P. Kerger, D. E. Bernal Neira, Z. Gonzalez Izquierdo, E. G. Rieffel, Mind the  $\bar{O}$ : Asymptotically Better, but Still Impractical, Quantum Distributed Algorithms, *Algorithms* 16 (2023). URL: <https://www.mdpi.com/1999-4893/16/7/332>. doi:10.3390/a16070332.
- [51] S. Hadfield, On the representation of Boolean and real functions as Hamiltonians for quantum computing, *ACM Transactions on Quantum Computing* 2 (2021) 1–21. doi:10.1145/3478519.
- [52] N. P. Sawaya, A. T. Schmitz, S. Hadfield, Encoding trade-offs and design toolkits in quantum algorithms for discrete optimization: coloring, routing, scheduling, and other problems, *Quantum* 7 (2023) 1111. doi:10.22331/q-2023-09-14-1111.
- [53] J. van de Wetering, ZX-calculus for the working quantum computer scientist, [arXiv preprint arXiv:2012.13966](https://arxiv.org/abs/2012.13966) (2020). [arXiv:2012.13966](https://arxiv.org/abs/2012.13966).
- [54] T. Stollenwerk, S. Hadfield, Diagrammatic Analysis for Parameterized Quantum Circuits, in: *Proceedings 19th International Conference on Quantum Physics and Logic (QPL 2022)*, Wolfson College, Oxford, UK, 27 June - 1 July 2022, volume 394, *Electronic Proceedings in Theoretical Computer Science*, 2023, pp. 262–301. URL: <http://cgi.cse.unsw.edu.au/~eptcs/paper.cgi?QPL2022.15>. doi:10.4204/EPTCS.394.15.
- [55] C. Zhao, X.-S. Gao, Analyzing the barren plateau phenomenon in training quantum neural networks with the ZX-calculus, *Quantum* 5 (2021) 466. doi:10.22331/q-2021-06-04-466.
- [56] A. Toumi, R. Yeung, G. de Felice, Diagrammatic differentiation for quantum machine learning, 2021. [arXiv:2103.07960](https://arxiv.org/abs/2103.07960).
- [57] S. Mandrà, A. Akbari Asanjan, L. Brady, A. Lott, D. E. Bernal Neira, PySA: Fast Simulated Annealing in Native Python, 2023. URL: <https://github.com/nasa/pysa>.
- [58] P. Maciel Xavier, P. Ripper, T. Andrade, J. Dias Garcia, N. Maculan, D. E. Bernal Neira, QUBO.jl: A Julia Ecosystem for Quadratic Unconstrained Binary Optimization, [arXiv e-prints](https://arxiv.org/abs/2307.02577) (2023) [arXiv:2307.02577](https://arxiv.org/abs/2307.02577). doi:10.48550/arXiv.2307.02577. [arXiv:2307.02577](https://arxiv.org/abs/2307.02577).
- [59] R. Brown, D. E. Bernal Neira, D. Venturelli, M. Pavone, A Copsitive Framework for Analysis of Hybrid Ising-Classical Algorithms, 2023. [arXiv:2207.13630](https://arxiv.org/abs/2207.13630).



- [60] R. Brown, et al., Accelerating Continuous Variable Coherent Ising Machines via Momentum, 2024. [arXiv:2401.12135](https://arxiv.org/abs/2401.12135).
- [61] M. Kim, S. Mandrà, D. Venturelli, K. Jamieson, Physics-inspired heuristics for soft MIMO detection in 5G new radio and beyond, in: Proceedings of the 27th Annual International Conference on Mobile Computing and Networking, MobiCom '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 42–55. URL: <https://doi.org/10.1145/3447993.3448619>. doi:10.1145/3447993.3448619.
- [62] N. Mohseni, et al., Ising machines as hardware solvers of combinatorial optimization problems, *Nature Reviews Physics* 4 (2022) 363–379.
- [63] A. K. Singh, K. Jamieson, D. Venturelli, P. McMahon, Ising Machines' Dynamics and Regularization for Near-Optimal Large and Massive MIMO Detection, 2021. [arXiv:2105.10535](https://arxiv.org/abs/2105.10535).
- [64] A. Taassob, et al., Neural Deep Operator Networks representation of Coherent Ising Machine Dynamics, *Neurips 2023 Workshop Machine Learning with New Compute Paradigms* (2023).
- [65] E. Witten, A new look at the path integral of quantum mechanics, *Surv. Differ. Geom.* 15 (2010) 345–421. URL: <https://arxiv.org/abs/1009.6032>.
- [66] E. Witten, Analytic continuation of Chern-Simons theory, *AMS/IP Stud. Adv. Math* 50 (2011) 347. URL: <https://arxiv.org/abs/1001.2933>.
- [67] T. C. Mooney, J. Bringewatt, N. C. Warrington, L. T. Brady, Lefschetz thimble quantum Monte Carlo for spin systems, *Phys. Rev. B* 106 (2022) 214416. URL: <https://link.aps.org/doi/10.1103/PhysRevB.106.214416>. doi:10.1103/PhysRevB.106.214416.
- [68] G. Bendall, et al., Greedy-type resistance of combinatorial problems, *Discrete Optimization* 3 (2006) 288–298.
- [69] A. Möbius, et al., Optimization by Thermal Cycling, *PRL* 79 (1997) 4297–4301.
- [70] A. Barzegar, A. Kankani, S. Mandrà, H. G. Katzgraber, Optimization and benchmarking of the thermal cycling algorithm, *Phys. Rev. E* 104 (2021) 035302. URL: <https://link.aps.org/doi/10.1103/PhysRevE.104.035302>. doi:10.1103/PhysRevE.104.035302.
- [71] Z. Zhu, et al., Efficient cluster algorithm for spin glasses in any space dimension, *PRL* 115 (2015) 077201.
- [72] S. Mandra, et al., Strengths and weaknesses of weak-strong cluster problems: A detailed overview of state-of-the-art classical heuristics versus quantum approaches, *PRA* 94 (2016) 022337.
- [73] S. Mandra, et al., A deceptive step towards quantum speedup detection, *Quantum Science and Technology* 3 (2018) 04LT01.
- [74] N. Gao, M. Wilson, T. Vandal, W. Vinci, R. Nemani, E. Rieffel, High-Dimensional Similarity Search with Quantum-Assisted Variational Autoencoder, in: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 956–964. URL: <https://doi.org/10.1145/3394486.3403138>. doi:10.1145/3394486.3403138.
- [75] A. Akbari Asanjan, M. Memarzadeh, P. A. Lott, E. Rieffel, S. Grabbe, Probabilistic Wildfire Segmentation Using Supervised Deep Generative Model from Satellite Imagery, *Remote Sensing* 15 (2023). URL: <https://www.mdpi.com/2072-4292/15/11/2718>. doi:10.3390/rs15112718.
- [76] M. Wilson, T. Vandal, T. Hogg, E. G. Rieffel, Quantum-assisted associative adversarial network: applying quantum annealing in deep learning, *Quantum Machine Intelligence* 3 (2021) 19. URL: <https://doi.org/10.1007/s42484-021-00047-9>. doi:10.1007/s42484-021-00047-9.
- [77] M. Wilson, R. Stromswold, F. Wudarski, S. Hadfield, N. M. Tubman, E. G. Rieffel, Optimizing quantum heuristics with meta-learning, *Quantum Machine Intelligence* 3 (2021) 13. URL: <https://doi.org/10.1007/s42484-020-00022-w>. doi:10.1007/s42484-020-00022-w.
- [78] D. O'Connor, W. Vinci, RBM-Flow and D-Flow: Invertible Flows with Discrete Energy Base Spaces, 2021. [arXiv:2012.13196](https://arxiv.org/abs/2012.13196).
- [79] T. Templin, M. Memarzadeh, W. Vinci, P. A. Lott, A. A. Asanjan, A. A. Armenakas, E. Rieffel, Anomaly Detection in Aeronautics Data with Quantum-compatible Discrete Deep Generative Model, 2023. [arXiv:2303.12302](https://arxiv.org/abs/2303.12302).
- [80] A. Kitaev, Anyons in an exactly solved model and beyond, *Annals of Physics* 321 (2006) 2–111. URL: <https://www.sciencedirect.com/science/article/pii/S0003491605002381>. doi:<https://doi.org/10.1016/j.aop.2005.10.005>, January Special Issue.
- [81] A. C. Y. Li, M. S. Alam, T. Iadecola, A. Jahin, J. Job, D. M. Kurkcuoglu, R. Li, P. P. Orth, A. B. Ozguler, G. N. Perdue, N. M. Tubman, Benchmarking variational quantum eigensolvers for the square-octagon-lattice Kitaev model, *Phys. Rev. Res.* 5 (2023) 033071. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.5.033071>. doi:10.1103/PhysRevResearch.5.033071.
- [82] D. Wecker, M. B. Hastings, M. Troyer, Progress towards practical quantum variational algorithms, *Phys. Rev. A* 92 (2015) 042303. URL: <https://link.aps.org/doi/10.1103/PhysRevA.92.042303>. doi:10.1103/PhysRevA.92.042303.
- [83] A. Jahin, A. C. Y. Li, T. Iadecola, P. P. Orth, G. N. Perdue, A. Macridin, M. S. Alam, N. M. Tubman, Fermionic approach to variational quantum simulation of Kitaev spin models, *Phys. Rev. A* 106 (2022) 022434. URL: <https://link.aps.org/doi/10.1103/PhysRevA.106.022434>. doi:10.1103/PhysRevA.106.022434.
- [84] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, et al., Observation of separated dynamics of charge and spin in the Fermi-Hubbard model, 2020. [arXiv:2010.07965](https://arxiv.org/abs/2010.07965).
- [85] R. Levy, Z. G. Izquierdo, Z. Wang, J. Marshall, J. Barreto, L. Fry-Bouriaux, D. T. O'Connor, P. A. Warburton, N. Wiebe, E. Rieffel, F. A. Wudarski, Towards solving the Fermi-Hubbard model via tailored quantum annealers, 2022. [arXiv:2207.14374](https://arxiv.org/abs/2207.14374).
- [86] V. Kremenetski, C. Mejuto-Zaera, S. J. Cotton, N. M. Tubman, Simulation of adiabatic quantum computing for molecular ground states, *The Journal of Chemical Physics* 155 (2021) 234106. URL: <https://doi.org/10.1063/5.0060124>. doi:10.1063/5.0060124.
- [87] A. J. C. Woitzik, P. K. Barkoutsos, F. Wudarski, A. Buchleitner, I. Tavernelli, Entanglement production and convergence properties of the variational quantum eigensolver, *Physical Review A* 102 (2020). URL: <http://dx.doi.org/10.1103/PhysRevA.102.042402>. doi:10.1103/physreva.102.042402.
- [88] A. Khan, B. K. Clark, N. M. Tubman, Pre-optimizing variational quantum eigensolvers with tensor networks, [arXiv:2310.12965](https://arxiv.org/abs/2310.12965) (2023).
- [89] 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, N. M. Tubman, Real time evolution for ultracompact Hamiltonian eigenstates on quantum hardware, 2021. [arXiv:2103.08563](https://arxiv.org/abs/2103.08563).
- [90] E. J. Gustafson, A. C. Y. Li, A. Khan, J. Kim, D. M. Kurkcuoglu, M. S. Alam, P. P. Orth, A. Rahmani, T. Iadecola, Preparing quantum many-body scar states on quantum computers, *Quantum* 7 (2023) 1171. URL: <http://dx.doi.org/10.22331/q-2023-11-07-1171>. doi:10.22331/q-2023-11-07-1171.
- [91] D. B. Chamaki, S. Hadfield, K. Klymko, B. O'Gorman, N. M. Tubman, Self-consistent Quantum Iteratively Sparsified Hamiltonian method (SQiSH): A new algorithm for efficient Hamiltonian simulation and compression, 2022. [arXiv:2211.16522](https://arxiv.org/abs/2211.16522).
- [92] O. A. von Lilienfeld, K. Burke, Retrospective on a decade of machine learning for chemical discovery, *Nature Communications* 11 (2020) 4895. URL: <https://doi.org/10.1038/s41467-020-18556-9>. doi:10.1038/s41467-020-18556-9.
- [93] D. Pfau, J. S. Spencer, A. G. D. G. Matthews, W. M. C. Foulkes, Ab initio solution of the many-electron Schrödinger equation with deep neural networks, *Physical Review Research* 2 (2020). URL: <http://dx.doi.org/10.1103/PhysRevResearch.2.033429>. doi:10.1103/physrevresearch.2.033429.
- [94] M. Wilson, N. Gao, F. Wudarski, E. Rieffel, N. M. Tubman, Simulations of state-of-the-art fermionic neural network wave functions with diffusion Monte Carlo, 2021. [arXiv:2103.12570](https://arxiv.org/abs/2103.12570).
- [95] M. Wilson, S. Moroni, M. Holzmann, N. Gao, F. Wudarski, T. Vegge, A. Bhowmik, Neural network ansatz for periodic wave functions and the homogeneous electron gas, *Physical Review B* 107 (2023). URL: <http://dx.doi.org/10.1103/PhysRevB.107.235139>. doi:10.1103/physrevb.107.235139.
- [96] D. E. Bernal, A. Ajagekar, S. M. Harwood, S. T. Stober, D. Trenev, F. You, Perspectives of quantum computing for chemical engineering, *AIChE Journal* 68 (2022) e17651. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/aic.17651>. doi:<https://doi.org/10.1002/aic.17651>.
- [97] D. E. Bernal Neira, C. D. Laird, S. M. Harwood, D. Trenev, D. Venturelli, Impact of Emerging Computing Architectures and Opportuni-



- ties for Process Systems Engineering Applications, in: Foundations of Computer-Aided Process Operations (FOCAPO)/Chemical Process Control (CPC) 2023, 2023.
- [98] S. J. Cotton, A truncated Davidson method for the efficient “chemically accurate” calculation of full configuration interaction wavefunctions without any large matrix diagonalization, *The Journal of Chemical Physics* 157 (2022) 224105. doi:10.1063/5.0115796.
- [99] N. Suri, J. Barreto, S. Hadfield, N. Wiebe, F. Wudarski, J. Marshall, Two-Unitary Decomposition Algorithm and Open Quantum System Simulation, *Quantum* 7 (2023) 1002. URL: <https://doi.org/10.22331/q-2023-05-15-1002>. doi:10.22331/q-2023-05-15-1002.
- [100] D. M. Kurkcuoglu, M. S. Alam, J. A. Job, A. C. Y. Li, A. Macridin, G. N. Perdue, S. Providence, Quantum simulation of  $\phi^4$  theories in qudit systems, 2022. arXiv:2108.13357.
- [101] H. Lamm, S. Lawrence, Y. Yamauchi (NuQS), General Methods for Digital Quantum Simulation of Gauge Theories, *Phys. Rev. D* 100 (2019) 034518. doi:10.1103/PhysRevD.100.034518. arXiv:1903.08807.
- [102] M. S. Alam, S. Hadfield, H. Lamm, A. C. Li, Primitive quantum gates for dihedral gauge theories, *Physical Review D* 105 (2022). URL: <http://dx.doi.org/10.1103/PhysRevD.105.114501>. doi:10.1103/physrevd.105.114501.
- [103] C. Charles, E. J. Gustafson, E. Hardt, F. Herren, N. Hogan, H. Lamm, S. Starecheski, R. S. V. de Water, M. L. Wagman, Simulating  $\mathbb{Z}_2$  lattice gauge theory on a quantum computer, 2023. arXiv:2305.02361.
- [104] E. J. Gustafson, H. Lamm, F. Lovelace, Primitive Quantum Gates for an  $SU(2)$  Discrete Subgroup: Binary Octahedral (2023). arXiv:2312.10285.
- [105] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, et al., Quantum Supremacy Using a Programmable Superconducting Processor, *Nature* 574 (2019) 505–510. URL: <https://www.nature.com/articles/s41586-019-1666-5>. doi:10.1038/s41586-019-1666-5.
- [106] B. Villalonga, D. Lyakh, S. Boixo, H. Neven, T. S. Humble, R. Biswas, E. G. Rieffel, A. Ho, S. Mandrà, Establishing the quantum supremacy frontier with a 281 Pflop/s simulation, *Quantum Science and Technology* 5 (2020) 034003. URL: <https://dx.doi.org/10.1088/2058-9565/ab7eeb>. doi:10.1088/2058-9565/ab7eeb.
- [107] D. Gottesman, Fault-tolerant quantum computation with higher-dimensional systems, in: NASA International Conf. on Quantum Computing and Quantum Communications, Springer, 1998, pp. 302–313.
- [108] S. Aaronson, D. Gottesman, Improved simulation of stabilizer circuits, *PRA* 70 (2004) 052328.
- [109] S. Bravyi, D. Gosset, Improved classical simulation of quantum circuits dominated by Clifford gates, *PRL* 116 (2016) 250501.
- [110] X. Mi, P. Roushan, C. Quintana, S. Mandrà, J. Marshall, et al., Information scrambling in quantum circuits, *Science* 374 (2021) 1479–1483. URL: <https://www.science.org/doi/abs/10.1126/science.abg5029>. doi:10.1126/science.abg5029.
- [111] D. Aharonov, X. Gao, Z. Landau, Y. Liu, U. Vazirani, A polynomial-time classical algorithm for noisily random circuit sampling, in: Proc. of the 55th Annual ACM Symposium on Theory of Computing, 2023, pp. 945–957.
- [112] I. L. Markov, Y. Shi, Simulating quantum computation by contracting tensor networks, *SIAM Journal on Computing* 38 (2008) 963–981.
- [113] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, et al., Phase transition in Random Circuit Sampling, 2023. arXiv:2304.11119.
- [114] K. Kechedzhi, S. Isakov, S. Mandrà, B. Villalonga, X. Mi, S. Boixo, V. Smelyanskiy, Effective quantum volume, fidelity and computational cost of noisy quantum processing experiments, *FGCS* 153 (2024) 431–441.
- [115] B. Bertini, P. Kos, T. c. v. Prosen, Entanglement Spreading in a Minimal Model of Maximal Many-Body Quantum Chaos, *Phys. Rev. X* 9 (2019) 021033. URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.021033>. doi:10.1103/PhysRevX.9.021033.
- [116] J. Gray, S. Kourtis, Hyper-optimized tensor network contraction, *Quantum* 5 (2021) 410.
- [117] F. Pan, K. Chen, P. Zhang, Solving the Sampling Problem of the Sycamore Quantum Circuits, *PRL* 129 (2022) 090502.
- [118] S. Arnborg, D. G. Corneil, A. Proskurowski, Complexity of finding embeddings in  $ak$ -tree, *SIAM Journal on Algebraic Discrete Methods* 8 (1987) 277–284.
- [119] F. Verstraete, V. Murg, J. I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Advances in physics* 57 (2008) 143–224. doi:<https://doi.org/10.1080/14789940801912366>.
- [120] G. Vidal, Class of quantum many-body states that can be efficiently simulated, *PRL* 101 (2008) 110501. doi:10.1103/PhysRevLett.101.110501.
- [121] J. Marshall, D. Kafri, Incoherent Approximation of Leakage in Quantum Error Correction, 2023. URL: <http://arxiv.org/abs/2312.10277>. doi:10.48550/arXiv.2312.10277. arXiv:2312.10277.
- [122] J. J. Renema, A. Menssen, W. R. Clements, G. Triginer, W. S. Kolthammer, I. A. Walmsley, Efficient Classical Algorithm for Boson Sampling with Partially Distinguishable Photons, *Phys. Rev. Lett.* 120 (2018) 220502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.220502>. doi:10.1103/PhysRevLett.120.220502.
- [123] C. Sparrow, Quantum Interference in Universal Linear Optical Devices for Quantum Computation and Simulation (2017). URL: <http://spiral.imperial.ac.uk/handle/10044/1/67638>. doi:10.25560/67638.
- [124] J. Marshall, Distillation of Indistinguishable Photons, *Phys. Rev. Lett.* 129 (2022) 213601. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.213601>. doi:10.1103/PhysRevLett.129.213601.
- [125] J. Saied, J. Marshall, N. Anand, E. G. Rieffel, General protocols for the efficient distillation of indistinguishable photons, arXiv preprint arXiv:2404.14217 (2024). doi:10.48550/arXiv.2404.14217.
- [126] F. Somhorst, B. K. Sauër, S. van den Hoven, J. J. Renema, Photon distillation schemes with reduced resource costs based on multiphoton Fourier interference, arXiv preprint arXiv:2404.14262 (2024). doi:10.48550/arXiv.2404.14262.
- [127] S. Rahimi-Keshari, T. C. Ralph, C. M. Caves, Sufficient Conditions for Efficient Classical Simulation of Quantum Optics, *Physical Review X* 6 (2016) 021039. URL: <https://link.aps.org/doi/10.1103/PhysRevX.6.021039>. doi:10.1103/PhysRevX.6.021039.
- [128] J. Marshall, N. Anand, Simulation of Quantum Optics by Coherent State Decomposition, *Optica Quantum* 1 (2023) 78–93. URL: <https://opg.optica.org/opticaq/abstract.cfm?uri=opticaq-1-2-78>. doi:10.1364/OPTICAQ.504311.
- [129] A. A. Mele, Introduction to Haar Measure Tools in Quantum Information: A Beginner’s Tutorial, *Quantum* 8 (2024) 1340. URL: <https://doi.org/10.22331/q-2024-05-08-1340>. doi:10.22331/q-2024-05-08-1340.
- [130] C. Dankert, R. Cleve, J. Emerson, E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, *Phys. Rev. A* 80 (2009) 012304. URL: <https://link.aps.org/doi/10.1103/PhysRevA.80.012304>. doi:10.1103/PhysRevA.80.012304.
- [131] H.-Y. Huang, R. Kueng, J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nature Physics* 16 (2020) 1050–1057. URL: <http://dx.doi.org/10.1038/s41567-020-0932-7>. doi:10.1038/s41567-020-0932-7.
- [132] J. Saied, J. Marshall, N. Anand, S. Grabbe, E. G. Rieffel, Advancing quantum networking: some tools and protocols for ideal and noisy photonic systems, in: P. R. Hemmer, A. L. Migdall (Eds.), *Quantum Computing, Communication, and Simulation IV*, SPIE, 2024. URL: <http://dx.doi.org/10.1117/12.3023843>. doi:10.1117/12.3023843.
- [133] N. P. Sawaya, D. Marti-Dafcik, Y. Ho, D. P. Tabor, D. E. Bernal Neira, A. B. Magann, S. Premaratne, P. Dubej, A. Matsuura, N. Bishop, et al., HamLib: A library of Hamiltonians for benchmarking quantum algorithms and hardware, in: 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), volume 2, IEEE, 2023, pp. 389–390. doi:10.1109/QCE57702.2023.10296.
- [134] T. Lubinski, C. Coffrin, C. McGeoch, P. Sathe, J. Apanavicius, D. E. Bernal Neira, Optimization Applications as Quantum Performance Benchmarks, 2023. arXiv:2302.02278.
- [135] D. E. Bernal Neira, et al., Benchmarking the Operation of Quantum Heuristics and Ising Machines: Scoring Parameter Setting Strategies on Optimization Applications, arXiv:2402.10255 (2024).
- [136] D. E. Bernal Neira, R. Brown, P. Sathe, D. Venturelli, Stochastic Benchmark: toolkit for performance evaluation and parameter tuning

- of stochastic parameterized stochastic optimization solvers, 2023. URL: <https://github.com/usra-riacs/stochastic-benchmark>.
- [137] F. Wudarski, J. Marshall, A. Petukhov, E. Rieffel, Augmented fidelities for single-qubit gates, *Phys. Rev. A* 102 (2020) 052612. URL: <https://link.aps.org/doi/10.1103/PhysRevA.102.052612>. doi:10.1103/PhysRevA.102.052612.
- [138] J. Sud, J. Marshall, Z. Wang, E. Rieffel, F. A. Wudarski, Dual-map framework for noise characterization of quantum computers, *Phys. Rev. A* 106 (2022) 012606. URL: <https://link.aps.org/doi/10.1103/PhysRevA.106.012606>. doi:10.1103/PhysRevA.106.012606.
- [139] J. Marshall, F. Wudarski, S. Hadfield, T. Hogg, Characterizing local noise in QAOA circuits, *IOP SciNotes* 1 (2020) 025208. URL: <https://dx.doi.org/10.1088/2633-1357/abb0d7>. doi:10.1088/2633-1357/abb0d7.
- [140] F. Wudarski, Y. Zhang, M. I. Dykman, Nonergodic Measurements of Qubit Frequency Noise, *Phys. Rev. Lett.* 131 (2023) 230201. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.230201>. doi:10.1103/PhysRevLett.131.230201.
- [141] F. Wudarski, Y. Zhang, A. N. Korotkov, A. Petukhov, M. Dykman, Characterizing Low-Frequency Qubit Noise, *Phys. Rev. Appl.* 19 (2023) 064066. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.19.064066>. doi:10.1103/PhysRevApplied.19.064066.
- [142] T. McCourt, C. Neill, K. Lee, C. Quintana, Y. Chen, J. Kelly, J. Marshall, V. N. Smelyanskiy, M. I. Dykman, A. Korotkov, I. L. Chuang, A. G. Petukhov, Learning noise via dynamical decoupling of entangled qubits, *Phys. Rev. A* 107 (2023) 052610. URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.052610>. doi:10.1103/PhysRevA.107.052610.
- [143] B. Evert, Z. G. Izquierdo, J. Sud, H.-Y. Hu, S. Grabbe, E. G. Rieffel, M. J. Reagor, Z. Wang, Syncopated Dynamical Decoupling for Suppressing Crosstalk in Quantum Circuits, 2024. [arXiv:2403.07836](https://arxiv.org/abs/2403.07836).
- [144] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, et al., Suppressing quantum errors by scaling a surface code logical qubit, *Nature* 614 (2023) 676–681. URL: <https://doi.org/10.1038/s41586-022-05434-1>. doi:10.1038/s41586-022-05434-1.
- [145] S. Alam, et al., Dynamical Logical Qubits in the Bacon-Shor Code, [arXiv](https://arxiv.org/abs/2024) (2024).
- [146] H.-Y. Hu, R. LaRose, Y.-Z. You, E. Rieffel, Z. Wang, Logical shadow tomography: Efficient estimation of error-mitigated observables, 2022. [arXiv:2203.07263](https://arxiv.org/abs/2203.07263).
- [147] A. A. Akhtar, N. Anand, J. Marshall, Y.-Z. You, Dual-Unitary Classical Shadow Tomography, 2024. [arXiv:2404.01068](https://arxiv.org/abs/2404.01068).
- [148] N. Suri, J. Saied, D. Venturelli, Uniformly Decaying Subspaces for Error Mitigated Quantum Computation, 2024. [arXiv:2403.00163](https://arxiv.org/abs/2403.00163).
- [149] X. Fei, L. T. Brady, J. Larson, S. Leyffer, S. Shen, Binary Control Pulse Optimization for Quantum Systems, *Quantum* 7 (2023) 892. URL: <https://doi.org/10.22331/q-2023-01-04-892>. doi:10.22331/q-2023-01-04-892.
- [150] X. Fei, L. T. Brady, J. Larson, S. Leyffer, S. Shen, Switching Time Optimization for Binary Quantum Optimal Control, 2023. [arXiv:2308.03132](https://arxiv.org/abs/2308.03132).
- [151] X. Fei, L. T. Brady, J. Larson, S. Leyffer, S. Shen, Binary Quantum Control Optimization with Uncertain Hamiltonians, 2024. [arXiv:2401.10120](https://arxiv.org/abs/2401.10120).
- [152] A. B. Özgüler, D. Venturelli, Numerical gate synthesis for quantum heuristics on bosonic quantum processors, *Frontiers in Physics* (2022) 724.
- [153] M. S. Alam, S. Belomestnykh, N. Bornman, G. Cancelo, Y.-C. Chao, M. Checchin, V. S. Dinh, A. Grassellino, E. J. Gustafson, R. Harnik, C. R. H. McRae, Z. Huang, K. Kapoor, T. Kim, J. B. Kowalkowski, M. J. Kramer, Y. Krasnikova, P. Kumar, D. M. Kurcuoglu, H. Lamm, A. L. Lyon, D. Milathianaki, A. Murthy, J. Mutus, I. Nekrashevich, J. Oh, A. B. Özgüler, G. N. Perdue, M. Reagor, A. Romanenko, J. A. Sauls, L. Stefanazzi, N. M. Tubman, D. Venturelli, C. Wang, X. You, D. M. T. van Zanten, L. Zhou, S. Zhu, S. Zorzetti, Quantum computing hardware for HEP algorithms and sensing, 2022. [arXiv:2204.08605](https://arxiv.org/abs/2204.08605).
- [154] D. Xu, A. B. Özgüler, G. Di Guglielmo, N. Tran, G. N. Perdue, L. Carloni, F. Fahim, Neural network accelerator for quantum control, in: 2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS), IEEE, 2022, pp. 43–49. doi:10.1109/QCS56647.2022.00010.
- [155] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, et al., Evidence for the utility of quantum computing before fault tolerance, *Nature* 618 (2023) 500–505. doi:10.1038/s41586-023-06096-3.
- [156] M. Kim, S. Mandrà, D. Venturelli, K. Jamieson, Physics-inspired heuristics for soft MIMO detection in 5G new radio and beyond, in: Proceedings of the 27th Annual International Conference on Mobile Computing and Networking, 2021, pp. 42–55. doi:10.1145/3447993.3448619.
- [157] M. Aramon, G. Rosenberg, E. Valiante, T. Miyazawa, H. Tamura, H. G. Katzgraber, Physics-inspired optimization for quadratic unconstrained problems using a digital annealer, *Frontiers in Physics* 7 (2019) 48. doi:10.3389/fphy.2019.00048.
- [158] S. Mertens, M. Mézard, R. Zecchina, Threshold values of random K-SAT from the cavity method, *Random Structures & Algorithms* 28 (2006) 340–373. doi:10.1002/rsa.20090.
- [159] M. Mézard, G. Parisi, R. Zecchina, Analytic and algorithmic solution of random satisfiability problems, *Science* 297 (2002) 812–815. doi:10.1126/science.10732.
- [160] M. Mézard, R. Zecchina, Random k-satisfiability problem: From an analytic solution to an efficient algorithm, *Physical Review E* 66 (2002) 056126. doi:10.1103/PhysRevE.66.056126.
- [161] S. Kirkpatrick, B. Selman, Critical behavior in the satisfiability of random boolean expressions, *Science* 264 (1994) 1297–1301. doi:10.1126/science.264.5163.12.
- [162] G. Qu, M. Potkonjak, Hiding signatures in graph coloring solutions, in: International Workshop on Information Hiding, Springer, 1999, pp. 348–367.
- [163] W. Barthel, A. K. Hartmann, M. Leone, F. Ricci-Tersenghi, M. Weigt, R. Zecchina, Hiding solutions in random satisfiability problems: A statistical mechanics approach, *Physical review letters* 88 (2002) 188701. doi:10.1103/PhysRevLett.88.188701.
- [164] F. Krzakala, L. Zdeborová, Hiding quiet solutions in random constraint satisfaction problems, *Physical review letters* 102 (2009) 238701. doi:10.1103/PhysRevLett.102.238701.
- [165] L. Zdeborová, F. Krzakala, Quiet planting in the locked constraint satisfaction problems, *SIAM Journal on Discrete Mathematics* 25 (2011) 750–770. doi:10.1137/090750755.
- [166] F. Krzakala, M. Mézard, L. Zdeborová, Reweighted belief propagation and quiet planting for random k-sat, *Journal on Satisfiability, Boolean Modeling and Computation* 8 (2012) 149–171.
- [167] G. Sicuro, L. Zdeborová, The planted k-factor problem, *Journal of Physics A: Mathematical and Theoretical* 54 (2021) 175002. doi:10.1088/1751-8121/ab9e9d.
- [168] I. Hen, J. Job, T. Albash, T. F. Rønnow, M. Troyer, D. A. Lidar, Probing for quantum speedup in spin-glass problems with planted solutions, *Phys. Rev. A* 92 (2015) 042325. URL: <https://link.aps.org/doi/10.1103/PhysRevA.92.042325>. doi:10.1103/PhysRevA.92.042325.
- [169] J. Marshall, V. Martin-Mayor, I. Hen, Practical engineering of hard spin-glass instances, *Phys. Rev. A* 94 (2016) 012320. URL: <https://link.aps.org/doi/10.1103/PhysRevA.94.012320>. doi:10.1103/PhysRevA.94.012320.
- [170] S. Mandrà, G. Mossi, E. G. Rieffel, Generating Hard Ising Instances With Planted Solutions Using Post-Quantum Cryptographic Protocols, 2023. [arXiv:2308.09704](https://arxiv.org/abs/2308.09704).
- [171] I. Hen, E. G. Rieffel, M. Do, D. Venturelli, Phase Transitions in Planning Problems: Design and Analysis of Parameterized Families of Hard Planning Problems, in: 28th AIAA Conference on Artificial Intelligence, ARC-E-DAA-TN13195, 2014.
- [172] R. J. McEliece, A Public-Key Cryptosystem Based On Algebraic Coding Theory, *Deep Space Network Progress Report* 44 (1978) 114–116. URL: <https://api.semanticscholar.org/CorpusID:56502909>.
- [173] M. Mezard, A. Montanari, *Information, Physics, and Computation*, Oxford University Press, Inc., USA, 2009. doi:10.1093/acprof:oso/9780198570837.001.0001.
- [174] D. Perera, I. Akpabio, F. Hamze, S. Mandra, N. Rose, M. Aramon, H. G. Katzgraber, Chook – A comprehensive suite for generating binary opti-

- mization problems with planted solutions, 2021. [arXiv:2005.14344](https://arxiv.org/abs/2005.14344).
- [175] F. Hamze, J. Raymond, C. A. Pattison, K. Biswas, H. G. Katzgraber, Wishart planted ensemble: A tunably rugged pairwise Ising model with a first-order phase transition, *Physical Review E* 101 (2020) 052102. doi:10.1103/PhysRevE.101.052102.
- [176] I. Hen, Equation planting: a tool for benchmarking Ising machines, *Physical Review Applied* 12 (2019) 011003. doi:10.1103/PhysRevApplied.12.011003.
- [177] J. Claes, E. Rieffel, Z. Wang, Character Randomized Benchmarking for Non-Multiplicity-Free Groups With Applications to Subspace, Leakage, and Matchgate Randomized Benchmarking, *PRX Quantum* 2 (2021) 010351. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.010351>. doi:10.1103/PRXQuantum.2.010351.
- [178] H. Zhu, Multiqubit Clifford groups are unitary 3-designs, *Physical Review A* 96 (2017) 062336. doi:10.1103/PhysRevA.96.062336.
- [179] Z. Webb, The Clifford group forms a unitary 3-design, 2016. [arXiv:1510.02769](https://arxiv.org/abs/1510.02769).
- [180] R. Kueng, et al., Qubit stabilizer states are complex projective 3-designs, 2015. [arXiv:1510.02767](https://arxiv.org/abs/1510.02767).
- [181] N. Anand, et al., Qudit designs with applications to coherent states, 2024. In preparation.
- [182] J. Marshall, L. Campos Venuti, P. Zanardi, Noise Suppression via Generalized-Markovian Processes, *Physical Review A* 96 (2017) 052113. URL: <https://link.aps.org/doi/10.1103/PhysRevA.96.052113>. doi:10.1103/PhysRevA.96.052113.
- [183] E. Vlachos, H. Zhang, V. Maurya, J. Marshall, T. Albash, E. M. Levenson-Falk, Master equation emulation and coherence preservation with classical control of a superconducting qubit, *Phys. Rev. A* 106 (2022) 062620. URL: <https://link.aps.org/doi/10.1103/PhysRevA.106.062620>. doi:10.1103/PhysRevA.106.062620.
- [184] H. M. Wiseman, E. G. Cavalcanti, E. G. Rieffel, A "thoughtful" Local Friendliness no-go theorem: a prospective experiment with new assumptions to suit, *Quantum* 7 (2023) 1112. URL: <https://doi.org/10.22331/q-2023-09-14-1112>. doi:10.22331/q-2023-09-14-1112.
- [185] B. Swingle, Unscrambling the Physics of Out-of-Time-Order Correlators, *Nature Physics* 14 (2018) 988–990. doi:10.1038/s41567-018-0295-5.
- [186] S. Xu, B. Swingle, Scrambling Dynamics and Out-of-Time-Ordered Correlators in Quantum Many-Body Systems, *PRX Quantum* 5 (2024) 010201. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.5.010201>. doi:10.1103/PRXQuantum.5.010201.
- [187] A. Kitaev, A simple model of quantum holography (part 1), 2015. URL: <http://online.kitp.ucsb.edu/online/entangled15/kitaev/>.
- [188] A. Nahum, S. Vijay, J. Haah, Operator Spreading in Random Unitary Circuits, *Phys. Rev. X* 8 (2018) 021014. URL: <https://link.aps.org/doi/10.1103/PhysRevX.8.021014>. doi:10.1103/PhysRevX.8.021014.
- [189] P. Zanardi, Entanglement of quantum evolutions, *Phys. Rev. A* 63 (2001) 040304. URL: <https://link.aps.org/doi/10.1103/PhysRevA.63.040304>. doi:10.1103/PhysRevA.63.040304.
- [190] P. Hayden, J. Preskill, Black holes as mirrors: quantum information in random subsystems, *Journal of High Energy Physics* 2007 (2007) 120–120. URL: <http://dx.doi.org/10.1088/1126-6708/2007/09/120>. doi:10.1088/1126-6708/2007/09/120.
- [191] J. Dubail, Entanglement scaling of operators: a conformal field theory approach, with a glimpse of simulability of long-time dynamics in 1+1d, *Journal of Physics A: Mathematical and Theoretical* 50 (2017) 234001. URL: <http://dx.doi.org/10.1088/1751-8121/aa6f38>. doi:10.1088/1751-8121/aa6f38.
- [192] S. Bravyi, M. B. Hastings, F. Verstraete, Lieb-Robinson Bounds and the Generation of Correlations and Topological Quantum Order, *Phys. Rev. Lett.* 97 (2006) 050401. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.97.050401>. doi:10.1103/PhysRevLett.97.050401.
- [193] S. Mandrà, J. Marshall, E. G. Rieffel, R. Biswas, HybridQ: A Hybrid Simulator for Quantum Circuits, in: 2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS), 2021, pp. 99–109. doi:10.1109/QCS54837.2021.00015.
- [194] G. Styliaris, N. Anand, P. Zanardi, Information Scrambling over Bipartitions: Equilibration, Entropy Production, and Typicality, *Phys. Rev. Lett.* 126 (2021) 030601. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.030601>. doi:10.1103/PhysRevLett.126.030601.
- [195] P. Zanardi, N. Anand, Information scrambling and chaos in open quantum systems, *Phys. Rev. A* 103 (2021) 062214. URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.062214>. doi:10.1103/PhysRevA.103.062214.
- [196] N. Anand, P. Zanardi, BROTOCs and Quantum Information Scrambling at Finite Temperature, *Quantum* 6 (2022) 746. URL: <https://doi.org/10.22331/q-2022-06-27-746>. doi:10.22331/q-2022-06-27-746.
- [197] F. Andreadakis, N. Anand, P. Zanardi, Scrambling of algebras in open quantum systems, *Phys. Rev. A* 107 (2023) 042217. URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.042217>. doi:10.1103/PhysRevA.107.042217.
- [198] B. Barch, N. Anand, J. Marshall, E. Rieffel, P. Zanardi, Scrambling and operator entanglement in local non-hermitian quantum systems, *Phys. Rev. B* 108 (2023) 134305. URL: <https://link.aps.org/doi/10.1103/PhysRevB.108.134305>. doi:10.1103/PhysRevB.108.134305.
- [199] M. Marvian, D. A. Lidar, I. Hen, On the computational complexity of curing non-stoquastic Hamiltonians, *Nature Communications* 10 (2019). URL: <http://dx.doi.org/10.1038/s41467-019-09501-6>. doi:10.1038/s41467-019-09501-6.
- [200] J. Bringewatt, L. T. Brady, Simultaneous stoquasticity, *Phys. Rev. A* 105 (2022) 062601. URL: <https://link.aps.org/doi/10.1103/PhysRevA.105.062601>. doi:10.1103/PhysRevA.105.062601.
- [201] J. Marshall, D. Venturelli, I. Hen, E. G. Rieffel, Power of Pausing: Advancing Understanding of Thermalization in Experimental Quantum Annealers, *Physical Review Applied* 11 (2019) 044083. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.11.044083>. doi:10.1103/PhysRevApplied.11.044083.
- [202] T. Albash, J. Marshall, Comparing Relaxation Mechanisms in Quantum and Classical Transverse-Field Annealing, *Phys. Rev. Appl.* 15 (2021) 014029. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.15.014029>. doi:10.1103/PhysRevApplied.15.014029.
- [203] E. Kapit, V. Oganesyan, Noise-tolerant quantum speedups in quantum annealing without fine tuning, *Quantum Science and Technology* 6 (2021) 025013. URL: <https://dx.doi.org/10.1088/2058-9565/abd59a>. doi:10.1088/2058-9565/abd59a.
- [204] G. Mossi, V. Oganesyan, E. Kapit, Embedding quantum optimization problems using AC driven quantum ferromagnets, 2023. [arXiv:arXiv:2306.10632](https://arxiv.org/abs/2306.10632).
- [205] G. Grattan, B. A. Barton, S. Feeney, G. Mossi, P. Patnaik, J. C. Sagar, L. D. Carr, V. Oganesyan, E. Kapit, Exponential acceleration of macroscopic quantum tunneling in a Floquet Ising model, 2023. [arXiv:arXiv:2311.17814](https://arxiv.org/abs/2311.17814).
- [206] N. Suri, C. Wang, B. M. Hunt, D. Xiao, Superlattice engineering of topology in massive Dirac fermions, *Phys. Rev. B* 108 (2023) 155409. URL: <https://link.aps.org/doi/10.1103/PhysRevB.108.155409>. doi:10.1103/PhysRevB.108.155409.
- [207] M. Do, Z. Wang, B. O’Gorman, D. Venturelli, E. Rieffel, J. Frank, Planning for Compilation of a Quantum Algorithm for Graph Coloring, 2020. [arXiv:2002.10917](https://arxiv.org/abs/2002.10917).
- [208] D. Venturelli, M. Do, E. Rieffel, J. Frank, Temporal Planning for Compilation of Quantum Approximate Optimization Circuits, in: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, 2017, pp. 4440–4446. URL: <https://doi.org/10.24963/ijcai.2017/620>. doi:10.24963/ijcai.2017/620.
- [209] M. S. Alam, N. F. Berthussen, P. P. Orth, Quantum logic gate synthesis as a Markov decision process, *npj Quantum Information* 9 (2023). URL: <https://doi.org/10.1038/s41534-023-00766-w>. doi:10.1038/s41534-023-00766-w.
- [210] A. Gold, J. Paquette, A. Stockklauser, M. J. Reagor, M. S. Alam, A. Bestwick, N. Didier, A. Nersisyan, F. Oruc, A. Razavi, B. Scharmann, E. A. Sete, B. Sur, D. Venturelli, C. J. Winkleblack, F. Wudarski, M. Harburn, C. Rigetti, Entanglement Across Separate Silicon Dies in a Modular Superconducting Qubit Device, 2021. [arXiv:2102.13293](https://arxiv.org/abs/2102.13293).