

White Paper: “Mathematical and Computational Challenges in Quantum Computing” (Long Program, Fall 2023, Institute for Pure and Applied Mathematics)

Authors: Ahmad M. Alkadri, Abhinav Anand, Andrew Baczewski, Brandon Barton, Joao Basso, M. Gabriela Boada G., Wendy Billings, Paul Cazeaux, Zhiyan Ding, Aviva Englander, Di Fang, Katerina Gratsea, Jakob Huhn, Michael A. Jones, Marius Junge, Tyler Kharazi, Spencer Lee, Jiaqi Leng, Lin Lin, Diyi Liu, Agnieszka Międlar, Jose Morales E., Scott Nie, Ojas Parekh, Soo-Jong Rey, Rahul Sarkar, Philipp Schleich, Lucas Tecot, Ruizhe Zhang, Shuchen Zhu

Table of Contents:

1. Executive Summary
2. Introduction
3. Input Model and Explicit State Preparation
4. Quantum Scientific Computing
5. Quantum Optimization
6. Quantum Learning
7. Open Quantum Systems
8. Noisy Intermediate-Scale Quantum (NISQ) Computation
9. Fault-Tolerant Quantum Computation
10. List of Acronyms

1. Executive Summary

This document serves as a summary of the research activities and outcomes of the Long Program, “Mathematical and Computational Challenges in Quantum Computing” which was held at the Institute of Pure and Applied Mathematics (IPAM) from September 11 to December 15, 2023. This program embraced the grand challenge in quantum information science: harness the weirdness of quantum mechanics to yield a computational advantage. Shor’s algorithm, featured by the Green Family Lecture Series delivered by Professor Peter Shor during the program, provides an early and (still) striking example of such an advantage. Practical applications of quantum computers also face several major challenges: they can require high ingenuity in preparing, manipulating and reading classical data, and existing quantum devices are prone to high error rates. This program employed deep tools from mathematics, inspired by physical realities, to find algorithmic solutions to challenging applications.

This document surveys key topics identified by program participants:

Input Model and Explicit State Preparation: Building an input model and preparing quantum state are prior steps for quantum scientific computing. The input model essentially builds the Hamiltonian of interest on a quantum computer and allows further oracle access. Along with the rapid development of quantum algorithms for scientific computing, it has become increasingly important to understand the end-to-end complexity and a clear estimate for the input model cost will be crucial. It is also an important task to develop input models with optimal complexity. Participants have made progress with respect to the input model for partial differential equations, quantum chemistry and random graphs as well as hardware-aware input model.

Quantum Scientific Computation: The rapid progress in quantum computing provides the scientific computing community with unprecedented opportunities. From a mathematical perspective, quantum algorithms perform a sequence of matrix-vector multiplications using only unitary matrices. However, many scientific computing tasks are not formulated as multiplications of unitary matrices. Over the past few decades, with particularly exciting progress in recent years, ingenious methods have been devised to adapt non-unitary operations and express them in terms of unitary operations. Participants have achieved significant progress along the lines of unbounded Hamiltonian simulation, linear differential equation solvers, and early fault tolerant quantum algorithms.

Quantum Optimization: Optimization drives ubiquitous applications across engineering and the sciences. This program sparked new collaborations between applied mathematicians, physicists, and computer scientists towards developing new kinds of classical and quantum algorithms for fundamental classical and quantum optimization problems such as linear regression, nonconvex optimization, boolean constraint satisfaction, and approximating many-body local Hamiltonians. The algorithms that participants developed or have envisioned break new ground by: delivering polynomial

and superpolynomial quantum advantages, leveraging quantum mechanics in novel ways, or generalizing classical optimization techniques for physical problems.

Quantum Learning: Machine learning (ML) has had a transformative impact across nearly every area of science and technology. There are two main ways that machine learning is used in the field of quantum computing. First, classical machine learning models are being used to learn quantum many-body systems. Second, learning using quantum resources has established the field of quantum machine learning. Variational quantum circuits have been proposed as models for quantum learning. Using quantum computers to achieve a quantum advantage in machine learning is a relatively new idea and thus has many open problems. The best way to optimize a parameterized circuit, and how to choose these quantum circuits is still an open problem and debated issue.

Open Quantum Systems: Recently, the importance of open quantum systems has emerged in experimental, theoretical and algorithmic research. They also play an important role in understanding noise, in particular in near-term quantum devices. Mathematical tools, numerical methods and experimental work have led to fundamental insights in ground state preparation, simulation and algorithms, and the connection to error correction.

Noisy Intermediate-Scale Quantum (NISQ) Computation: There is considerable interest in the capabilities of present- and near-term quantum devices where fault-tolerant quantum computation is not possible. In recent years significant progress has been made in developing quantum hardware, software, and algorithms to enable quantum computations that challenge classical devices. Quantum error mitigation provides a mechanism to reduce the effect of noise in the quantum processor, generally at the cost of additional circuit repetitions and/or classical post-processing.

Fault-tolerant Quantum Computation: Fault-tolerant quantum computation is required for a large majority of currently proposed quantum algorithms and will likely be necessary to achieve meaningful quantum advantage. Strategies for quantum error correction (QEC) encode logical qubits across many noisy qubits, but such schemes require hardware that achieves a certain fidelity threshold. Although full fault tolerance is currently out of reach, progress is being made in the design of good quantum codes and their implementation. There is increasing interest in work for the early fault-tolerant quantum computing (EFTQC) regime, where logical qubits and operations are available but only at a limited scale. Additional work estimating the total cost of running algorithms in a fault-tolerant setting will play an important role in directing further error correction and algorithmic work.

2. Introduction

Quantum information science is a multi-disciplinary challenge with broad applications to computing, communication, security, and quantum sensing tasks. According to Scott Aaronson, quantum computation must pass the very high bar of providing provable quantum advantages with a relatively small number of qubits on a noisy device in order

to demonstrate its near-term usefulness. The rapid advancement in quantum computing presents unparalleled opportunities for the scientific community. However, fully harnessing the potential of quantum computers to achieve quantum advantage over classical computers is a significant challenge. While it may be tempting to think that exponential quantum speedups can be achieved by using n qubits to encode 2^n bits of classical information, the reality is more subtle, since the quantum algorithm must interact with classical processing systems. Therefore, we need to not only design efficient quantum algorithms using a relatively small number of quantum gates, but also thoughtfully consider input-output models, and the specific requirements of quantum algorithms.

We may dissect the overall quantum cost into three main categories: input, output, and running costs. A quantum algorithm starts with a standard state, which is then transformed by a unitary matrix to prepare the input state. The input cost is the quantum gate complexity required to implement this unitary matrix. The output cost relates to the quantum measurement process, which is performed on one or several qubits often at the end of the algorithm. The output cost is given by the number of repetitions necessary to reach a target precision. The running cost refers to the expense incurred in executing the quantum algorithm a single time. A comprehensive "end-to-end" analysis of quantum advantage requires considering input, output, and running costs.

Assessing the performance of quantum algorithms and identifying potential quantum advantages raises considerable challenges. Rigorous proof of asymptotic quantum advantage over the best-possible classical algorithm, ideally superpolynomial, is a gold standard; however, such proofs are difficult to establish and may be currently beyond reach. An alternative is to compare against the best-known classical algorithm; however, such quantum advantages may vanish under classical algorithmic advances. Empirical assessment of quantum algorithms is a natural strategy in absence of rigorous proofs; however, this may suggest advantages that are not sustainable as problem instances grow in size or complexity. Moreover, as quantum information science (QIS) is an exemplar interdisciplinary field, methodologies or metrics for performance assessment may be highly domain sensitive. For example, quantum complexity theory seeks to provide insight into the power and limitations of quantum computing from a worst-case asymptotic perspective; however, such a notion may not be appropriate for computing quantities to a practical precision, such as chemical accuracy, for a specific molecule.

Research in quantum computation has to reflect and embrace the challenges imposed by the limitations of existing quantum devices. For example, fault tolerance and quantum error correction are omnipresent in quantum information theory and quantum computation. Over the past fifty years, classical computer gates have observed significant improvements resulting in miniscule errors per gate. In contrast, physical quantum gates are subject to broader sources of error stemming from their quantum-mechanical nature. Errors present a critical challenge for Noisy Intermediate-Scale Quantum (NISQ) devices. However, just as there exist classical error correction codes, which, for example, enable us to smoothly communicate over cell phones, there are quantum analogues called

quantum error-correction (QEC) codes that are known to be able to correct errors in quantum computation. Developing and engineering robust and scalable QEC is a paramount challenge for quantum computing.

Nearly optimal QEC rests on deep mathematical insights. Quantum computation also draws inspiration from interactions with other fields. For example, recent theoretical findings suggest that quantum field theory and gravity rely on surprising properties of quantum entanglement and codes. Fundamental problems from condensed matter theory and quantum chemistry present challenges that quantum computation seems well suited to address. Quantum computation has had a profound impact on computational complexity theory, with many previously studied models of computation exhibiting unexpected connections to quantum mechanics. A particularly striking example enabled a recent breakthrough disproving a longstanding conjecture in operator algebras (Connes' embedding problem) through non-local games and entanglement.

3. Input Model and Explicit State Preparation

Introduction: The input model serves as the vital link between quantum hardware, quantum error correction, and quantum algorithms, playing a foundational role in the initial stages of scientific computing by facilitating the creation of a Hamiltonian on a quantum computer. Quantum algorithms are typically formulated with the assumption of access to the input model, which comprises a SELECT oracle and a PREPARE oracle. The SELECT oracle encodes information about the location of nonzero elements in the Hamiltonian, while the PREPARE oracle precisely provides the values of these nonzero elements. The SELECT oracle's characteristics often depend on the nature of the problem, such as sparse matrices or quantum many-body Hamiltonians. Details about the problem are then handled by the PREPARE oracle. The construction of the input model involves various techniques aimed at building a complex Hamiltonian from simpler components. These techniques include Linear Combination of Unitaries (LCU), Quantum Random Access Memory (QRAM), and Hamiltonian Embedding.

Linear Combination of Unitaries (LCU): LCU methods were originally conceived to achieve optimal or near-optimal query complexities for Hamiltonian simulation. These early results paved the way for the broader class of block-encoding techniques, which embed and operate on a subnormalized matrix inside a unitary matrix. As the name suggests, LCU requires that this matrix can be rendered as a linear combination of unitaries for which the 1-norm of the coefficients provide the subnormalization constant. The Pauli basis is conceptually convenient for quantum computation, but other choices of unitary bases might have performance advantages. In order to realize an LCU embedding, a register of ancilla qubits is required with a number of qubits that is logarithmic in the number of terms in the LCU expansion. The block-encoded matrix is then accessed through PREPARE and SELECT oracles, which respectively load the square root of the coefficients of the LCU expansion as amplitudes in the ancilla register and apply the individual terms of the expansion on some target register, controlled on the state of the ancilla register. The product PREPARE and SELECT oracles can be

combined to form a product of reflection operators known as a Szegedy walk, which enables the powerful simulation method known as qubitization.

One of the primary drawbacks of the LCU approach to block encoding (BE) is the fact that the subnormalization constant determines the cost of working with this representation of the matrix. For a matrix with entries of mixed sign, the fact that this is determined by a 1-norm is especially damaging. Other challenges include the implementation of the SELECT and PREPARE oracles. While techniques like unary iteration facilitate the efficient implementation of SELECT given access to controlled versions of the constituent unitaries, the complexity of implementing PREPARE for an arbitrary matrix can be prohibitive. While unstructured matrices suffer from the cost of preparing arbitrary states on the ancilla register, specially structured matrices can see dramatic improvements in the gate complexity required for the implementation of PREPARE.

Quantum Random Access Memory (QRAM): Numerous quantum algorithms offer speedup over their classical counterparts. These algorithms typically require frequent access to classical data, which encodes the problem instance. QRAM, if implemented, would facilitate the efficient storage and retrieval of classical data in quantum computing. To realize practical quantum benefits, it is imperative to pinpoint a viable implementation for this data access scheme.

While QRAM plays a crucial role in the success of quantum algorithms, realizing its implementation poses significant challenges. The primary hurdle lies in ensuring fault-tolerant quantum circuits. Achieving a high-fidelity T-gate is resource intensive, making it imperative to reduce the T-gate count in circuit design without compromising fidelity. Moreover, the effectiveness of QRAM depends on its query time aligning with the total runtime of the algorithm. Significant overhead in data retrieval could undermine the potential quantum advantage of an algorithm which uses QRAM. Recent research demonstrates the error resilience of QRAM. Consequently, practical QRAM implementations should prioritize maintaining this error resilience, aiming for a high success rate, and optimizing factors such as T-gate count, query duration, and qubit consumption.

Hamiltonian embedding: Block encodings and QRAM are advanced quantum input models that require fully fault-tolerant quantum computers. Recently, a technique named Hamiltonian embedding has been proposed for simulating sparse Hamiltonians on near-term quantum devices. Hamiltonian embedding simulates a desired Hamiltonian evolution by embedding it into the evolution of a large and structured quantum system. The large embedding Hamiltonian is often represented as a sum of 1- and 2-local Pauli operators, which allows efficient simulation on both gate-based and analog quantum devices. Hamiltonian embedding can be regarded as a new quantum input model of sparse matrices that does not require fault tolerance. With Hamiltonian embedding, it is possible to implement quantum algorithms for scientific computing (e.g., non-unitary quantum dynamics) and optimization (e.g., Quantum Hamiltonian Descent) on near-term realizable quantum hardware.

Work done at IPAM:

1. Improved Block Encoding for general second quantized Hamiltonian: Recent strides in Quantum Random-Access Classical Memory (QRACM) prompt a fresh approach to block encoding second-quantized Hamiltonians. In this innovative framework, we adopt a novel encoding for creation and annihilation operators, moving beyond traditional Jordan-Wigner and Bravyi-Kitaev transformations. This updated encoding brings about a consistent acceleration concerning T-gates and success probability. Notably, when tackling intricate problems in nuclear and particle physics, the new approach demonstrates a remarkable 50-fold increase in algorithmic speed for a Hamiltonian including a three-body interacting term. Moreover, the QRACM introduces a reduction in T-gate count complexity relative to the qubit number n . By synergizing these advancements, we present an enhanced block encoding for general second-quantized Hamiltonians.

2. Block encoding for partial differential equations: The solution of partial differential equations (PDEs) has been proposed as a possible application of quantum computers. Given a finite dimensional matrix representation of the PDE, e.g. by finite difference approximations, block encoding is the standard input model for implementing the matrix on a quantum computer. However, many algorithms simply assume access to these block encodings without providing explicit cost estimates for the circuit. For finite difference approximations, the result is a large sparse matrix with regular structure. Using the method of linear combination of unitaries, with a family of unitaries called the clock and shift operators, during this program we have constructed explicit block encodings for finite difference representation of elliptic PDEs which commonly arise in scientific applications which scale independently of the number of grid points with non-trivial boundary conditions. This makes the cost analysis for the solution of PDEs on the quantum computer more precise for a broader class of problems and enables tighter analysis of the feasibility of quantum solvers for this class of problems.

3. Block encoding for random graphs: Several quantum algorithms claiming exponential speedups assume the BE of certain sparse matrices. However, the oracle steps in the BE are often very non-trivial to construct in practice. Thus, quantum advantage without end-to-end complexity analysis remains challenging, even for well-studied problems such as the glued trees problem. Recent work has provided explicit constructions for the adjacency matrices of some simple graphs including the cycle graph (circulant matrices), trees, and symmetric stochastic matrices. To provide constructions of relevant problems where quantum algorithms have proven advantages, in this program, we've begun work on a large project for the block encoding of graphs without well-defined structure. The results from this work have significant implications into understanding what problem structures are amenable to quantum advantage, with applications to graph theory and quantum walks on graphs.

4. Hamiltonian embedding: With the explicit construction of Hamiltonian embeddings of graph Laplacians and differential operators, we demonstrate several important quantum applications on near-term quantum devices (e.g., the IonQ and QuEra quantum computers), including quantum walk on glued trees graphs and the simulation of real-space Schrödinger equations.

Outlook: Several quantum algorithms, particularly in areas like quantum machine learning, heavily depend on accessing quantum data instead of classical data. However, the anticipated quantum speedup may not materialize if the resource cost and overhead associated with quantum data access are substantial. Surprisingly, there is limited literature on frameworks intricately designed for the non-trivial storage of quantum data. To our knowledge, an early design known as Quantum Random-Access Quantum Memory (QRAQM) has been coined for the explicit purpose of storing quantum data. Unfortunately, the T-gate count for this scheme is currently linear, presenting a potential hindrance to its efficiency. We believe that there is room for improvement, and efforts can be directed toward enhancing the T-gate count to further optimize the performance of the QRAQM framework.

4. Quantum Scientific Computing

Introduction: The recent development of quantum algorithms has significantly advanced the frontier of using quantum computers for performing a wide range of scientific computing tasks. This includes solving numerical linear algebra problems for very large matrices, such as solving linear systems, eigenvalue and singular value transformations, matrix function evaluation, trace estimation, topological data analysis, as well as solving certain high-dimensional linear and nonlinear differential equations, including for both quantum systems (Hamiltonian simulation) and other classical applications. This program has brought together leading experts across different disciplines, including experts in solving related tasks using classical computers that can potentially inspire the development of new quantum algorithms; discussed recent progress made in the development of quantum algorithms for scientific computation, and the advances in classical algorithms; paved the path towards identifying and overcoming challenging problems in science, engineering, and various industrial and technological applications.

Unbounded Hamiltonian simulation: Simulation of the quantum dynamics (also called Hamiltonian simulation) was the original motivation for quantum computers and remains as one of the most basic, fundamental, and important tasks in quantum computing. Recent years have witnessed tremendous progress in developing and analyzing quantum algorithms for Hamiltonian simulation of bounded operators. However, many scientific and engineering problems require the efficient treatment of unbounded operators for infinite-dimensional Hilbert spaces, which frequently arise due to the discretization of differential operators. Such applications include molecular dynamics, first-quantized electronic structure theory, quantum differential equation solver, quantum optimization, and quantum machine learning. These problems bring additional challenges to quantum algorithms whose cost typically depends on the norm of the Hamiltonian, which can be infinite. For unbounded operators, even after suitable discretization, the norm of the Hamiltonian can be very large, and it was found that, in the worst case, discretization errors that arise from the spatial grid can remove the exponential speedups typically afforded by quantum simulation.

There have been recent developments dedicated to addressing this challenge, which can be categorized into two mainstream approaches. One approach is to consider the specific or average-case scenario of a particular problem, as opposed to the worst case, often leading to an effective Hilbert space (or manifold) of much smaller dimension. This includes low-energy subspace estimations, vector norm analysis, average-case analysis for random inputs, and observable error bounds. The other approach involves utilizing the interaction picture formulation that transforms the (time-independent) Hamiltonian with a large norm into a time-dependent Hamiltonian with a much smaller norm but is highly oscillatory in time. This requires that quantum algorithms have a cost independent (or only weakly dependent) of the oscillatory behavior. Tools from mathematical analysis, such as analysis for PDEs, pseudodifferential operators, and semiclassical analysis, have been proven to be particularly useful in advancing the provable efficiency of quantum algorithms involving unbounded operators.

Linear systems and differential equations: The solution of differential equations often requires a large amount of classical resources to obtain accurate solutions. In many cases, the differential equations are represented as an ordinary differential equation (ODE) involving a finite dimensional matrix. Some examples where this is the case include numerical solution of partial differential equations (PDE) where the spatial components of the differential operator are discretized to form a finite dimensional matrix ODE. Several classical methods for time evolving ODEs have been translated into asymptotically efficient quantum algorithms. However, in many cases both the quantum and classical algorithms suffer from dependence on the matrix norm and/or condition number depending on the time stepping scheme of choice. Implicit methods, which involve solving a linear system at each time step, remain a relatively unexplored direction in quantum algorithms. However, these approaches offer the advantage of being able to take large time-steps.

For the class of elliptic PDEs, many classical methods exist to effectively and efficiently precondition the linear system. This, in part, is due to the fact that with periodic boundary conditions the elliptic operator is diagonalizable by Fourier transform; given a block encoding of the matrix on a quantum computer, this step can be performed exponentially faster using the quantum Fourier transform. Preconditioning on a quantum computer for other classes of PDEs, however, still remains a relatively unexplored avenue. In particular, explicit block encodings, resource estimates, and success probabilities for implementations of the most popular classes of preconditioners (such as those that arise from factorization methods, sparse approximate inverses, and multilevel paradigms) remain largely unknown. Due to these limitations, it is difficult to assess the outlook for quantum advantage in the solution of generic PDEs.

Non-Hermitian quantum dynamics: Non-Hermitian quantum dynamics simulates the dynamics associated with a non-Hermitian operator. This is an inherently non-unitary process and is a challenging task for quantum computers. A significant advancement in quantum algorithms over the past decade is the development of the quantum singular value transformation (QSVT). However, simulating non-Hermitian quantum dynamics

requires performing quantum eigenvalue transformations of a potentially non-normal matrix, and cannot be cast into the framework of singular value transformations. The leading approach is somewhat complex and perhaps counterintuitive. It starts by treating the problem as an ODE, then discretizing this ODE over time, and converting it into a large linear system of equations, which can be solved by quantum linear system algorithms (QLSA). The ODE can be solved using a traditional time-marching strategy, similar to that employed in standard numerical ODE solvers. However, direct implementation leads to an excessively high output cost due to diminishing success probability. Recently, a time-marching strategy that avoids this issue has been developed. This new time-marching strategy also achieves optimal query complexity with respect to the initial state preparation, which reduces the input cost.

Another recent advancement is that the simulation of non-unitary quantum dynamics can be greatly simplified as a linear combination of Hamiltonian simulation (LCHS) problems. The LCHS approach not only streamlines the simulation process, but also achieves optimal state preparation cost.

Early fault-tolerant quantum algorithms: The primary goals of algorithms in the early fault-tolerant quantum (EFTQC) era are as follows: **1)** Use only a short circuit depth and simple gates; **2)** Rely on a small number of ancilla qubits; **3)** Demonstrate robustness to noise. One typical example of EFTQC algorithm development appears in the area of quantum phase estimation, where the aim is to estimate the energies of a given Hamiltonian, assuming access to the Hamiltonian simulation. In the context of quantum phase estimation, algorithmic complexities are characterized by the maximal (which determines circuit depth) and total Hamiltonian simulation (which determines running time). In contrast to the focus in the development of fault-tolerant quantum phase estimation algorithms, where efforts emphasize reducing total Hamiltonian simulation, early fault-tolerant algorithms aim to minimize maximal Hamiltonian simulation time and the required number of ancilla qubits while still maintaining a small total Hamiltonian simulation.

Work done at IPAM: Different directions were explored by the program participants:

1. *Superconvergence of unbounded Hamiltonian simulation:* For the simulation of unbounded Hamiltonians in the interaction picture, we looked at the (super)convergence of the second-order Magnus series expansion. In previous works, it has been observed numerically and conjectured that the Magnus series expansion with second order truncation exhibits fourth order superconvergence. At IPAM, we set out to prove this fact. If one tries to do this naively, one runs into severe problems due to an explosion of the number of terms for which one needs to prove commutator estimates. In fact, we had major difficulty in carrying out the proof this way, as many terms were in fact unbounded operators on L^2 . The way we managed to make some headway into this proof is to introduce some tools from semiclassical microlocal analysis. A key commutator estimate with unbounded operators has been completely established using the semiclassical pseudodifferential operator machinery.

2. *Quantum implicit differential equations solvers:* During the IPAM long program, our

focus was on establishing a general quantum time-stepping theory for implicit short-time integrators in PDEs, aiming to minimize computational dependence on spatial grid size. We specifically investigated elliptic and parabolic PDEs where the preconditioner can be explicitly implemented through simple quantum primitives. Leveraging the fast-invertibility of the differential part via quantum Fourier Transform, we obtained results for the backward Euler method with a cost that depends logarithmically on the number of grid points. We have also developed a quantum circuit in terms of simple primitives that implements an arbitrary Runge-Kutta method, and we will continue to explore the development of these methods for other classes of differential equations, as well as preconditioners for higher-order implicit schemes.

3. *Non-Hermitian quantum dynamics*: For simulating the non-Hermitian quantum dynamics, we have developed a family of identities that express general linear non-unitary evolution operators as a linear combination of unitary evolution operators, each solving a Hamiltonian simulation problem. This formulation can exponentially enhance the accuracy of the original LCHS formula. For the first time, this approach enables quantum algorithms to solve linear differential equations with both optimal state preparation cost and near-optimal scaling in matrix queries on all parameters.

4. *Quantum phase estimation*: During the program, we developed a novel early fault-tolerant quantum multiple eigenvalue estimation algorithm that demands only one ancilla qubit and employs simple quantum circuits, making it suitable for the implementation on the early fault-tolerant quantum computer. Importantly, it can be proved to achieve (nearly-)optimal complexity scaling in terms of maximal/total Hamiltonian evolution time across all regimes, surpassing previous quantum phase estimation methods.

Outlook: The development of quantum algorithms for scientific computing connects various areas of mathematics and quantum information science, including numerical algebra, numerical analysis, complex analysis, PDEs, harmonic, and functional analysis, etc. This connection enhances estimates and helps to lead to novel quantum algorithms which prompt exploration of new mathematical problems. For simulating non-Hermitian dynamics, an immediate open question is whether we can develop efficient quantum algorithms for systems that are not exponentially stable but only asymptotically stable with transient growth behaviors. It would also be of significant interest to simulate non-Hermitian dynamics on physical quantum devices using state-of-the-art input models.

5. Quantum Optimization

Introduction: Optimization problems drive ubiquitous applications across engineering and the sciences. Considering optimization in a quantum context suggests four natural categories: **1)** Quantum algorithms for classical optimization problems, **2)** Classical algorithms for quantum optimization problems, **3)** Quantum algorithms for quantum optimization problems, and **4)** Quantum-inspired classical algorithms applied to classical optimization problems. Our efforts during this program primarily fall into Categories 1-3.

A prototypical quantum optimization problem entails estimating properties of low-energy states of a local many-body Hamiltonian. Such problems are a cornerstone of condensed matter physics and quantum chemistry. They may also be viewed as quantum generalizations of certain types of classical optimization problems, such as boolean constraint satisfaction problems (CSPs), since the latter may be cast as a local Hamiltonian problem on qubits that is diagonal in the computational basis.

Quantum algorithms for classical optimization problems: Classical optimization is a promising area to demonstrate quantum advantages. Numerous quantum algorithms have been developed to solve continuous and discrete optimization problems more efficiently than using classical computers. However, demonstrating end-to-end super-polynomial quantum speedups in classical optimization remains challenging.

Linear regression is a fundamental continuous optimization problem, with myriad scientific applications, including statistics, operations research, and machine learning. The development of efficient algorithms for large-scale linear regression, both in classical and quantum computing is a key research focus. Exponential quantum speedups are achievable in preparing quantum states that encode linear regression solutions, though this advantage is limited if classical output is needed or if the input matrix is ill-conditioned. In more complex constrained optimization problems like linear programming (LP) and semidefinite programming (SDP), numerous quantum algorithms have been proposed, offering polynomial speedups over classical methods. On the other hand, Quantum Hamiltonian Descent (QHD) is a recently developed approach to solve classical optimization problems by leveraging quantum mechanics in a unique way, without any direct classical analog. It is derived from quantizing dynamic systems that model the continuous-time limit of classical gradient-based algorithms. QHD, by leveraging quantum tunneling, navigates more effectively through nonconvex optimization landscapes. Moreover, QHD's applicability extends beyond standard circuit-based implementations to analog quantum simulators, as detailed in the "Hamiltonian embedding" part of Section 2, thus opening up the possibility for testing quantum optimization methods on large-scale empirical benchmarks.

In the field of discrete optimization, a wide range of problems can be expressed as CSPs, while designing efficient CSP solvers remains a challenging problem. A classical-quantum hybrid approach for CSPs is the Quantum Approximate Optimization Algorithm (QAOA), which uses a classical optimizer to train a parameterized quantum ansatz that prepares a quantum state encoding a good approximation to the CSP. The popularity of QAOA has prompted significant interest in quantum approximation advantages for discrete optimization problems. While no such example is currently known, recent work completed during this program produced the first superpolynomial quantum advantage in approximating a discrete optimization problem, directed Maxcut, albeit in the streaming model of computation where space, rather than time, is the resource of interest. This is the first example of a superpolynomial streaming advantage for a natural problem, resolving a 16-year-old open problem.

Classical algorithms for quantum optimization problems: Approaches for solving local Hamiltonian problems have historically been classical. While heuristics based on physical insights have been the traditional approach, a recent thread has explored producing algorithms with rigorous guarantees by employing hierarchies of semidefinite programs (SDPs) to solve relaxed problems that yield bounds on quantities of interest. SDP hierarchies have recently yielded rigorous approximations for QMA-hard local Hamiltonian problems as well as the first hardness of approximation results. Approximating local Hamiltonians is intimately related to the quantum PCP conjecture, an outstanding open problem.

While hierarchies of SDPs in commuting variables are common in optimization, the application of non-commutative SDP hierarchies to physics is in its infancy. Such hierarchies capture physical relationships between observables and can accommodate underlying symmetries. SDP hierarchies implicitly represent physical observables and wave functions, and the sizes of the SDPs are independent of the dimension of the underlying Hilbert space. SDPs at low levels of the hierarchy can be solved efficiently and provide increasingly better approximations, converging to optimality, as one moves up the hierarchy. Recent applications to quantum field theory (QFT) have remarkably yielded precise numerical estimates, excelling beyond standard perturbation expansions and renormalization-group approaches.

This program has sparked new collaborations among quantum field theorists, theoretical computer scientists, and applied mathematicians who are uniquely qualified to further develop applications of SDP methods to quantum many-body physics and quantum chemistry.

Quantum algorithms for quantum optimization problems: It is natural to imagine that quantum computing may be better suited to offer advantages on inherently quantum problems rather than classical ones. A canonical example is approximating Quantum Maxcut, a local Hamiltonian problem closely related to the well-studied quantum Heisenberg model and classical Maxcut CSP. By bridging classical and quantum techniques, several algorithms have recently been developed for this problem with increasingly better approximation guarantees. However, fundamental questions related to optimal approximability and hardness of approximation remain unanswered. Even less is known for the general local Hamiltonian problem.

Work done at IPAM: During the program, several projects emerged in the aforementioned categories. We proposed a new quantum linear regression solver with quadratic quantum speedup without depending on any data-related parameter. This addresses shortcomings of previous QLSA-based approaches that only worked for well-conditioned inputs. We made progress to quantitatively characterize the quantum advantage of QHD for nonconvex optimization, proving that QHD can find the global minimum of certain challenging nonconvex optimization problems in polynomial time, whereas most state-of-the-art classical optimizers exhibit super-polynomial scaling in their search for a good solution. We also managed to extend the idea of QHD to solve linear programming with provable end-to-end quantum speedups. We introduced a new

method called spectral folding quantum optimization, demonstrating that solutions within a constant fraction of optimal solutions is obtainable in polynomial time. This work clarified the mechanisms of hardness for the CSP hypergraph problem MAX-3-XORSAT, separating the hardness of exact and approximate quantum optimization. We also discussed a potential correspondence between local Hamiltonians and quantum spin glasses, which would lift the analogous classical result where a random CSP has an associated spin glass whose ground state energy encodes the best approximation ratio achievable for that CSP. If such correspondence holds, understanding of quantum spin glasses would lead to better understanding of local Hamiltonians.

Outlook: Given the extensive literature and broad interest within computational sciences, the interplay between quantum computing and optimization will continue to thrive as an active research field. The pursuit of strong theoretical guarantees for quantum optimization algorithms remains an important research direction, for example, super-Grover speedup for classical optimization problems such as CSP and even super-polynomial separations between quantum and classical for nonconvex problems. Combining classical algorithms with well-known approximation guarantees may provide one promising route to proving advantages in quantum optimization algorithms. On the other hand, considering the rapid advancement of quantum engineering and the intricate co-design between quantum algorithms and realistic physical hardware, achieving end-to-end quantum advantage for optimization problems emerging from application domains may remain a challenging endeavor in the near future.

6. Quantum Learning

Introduction: The intersection between machine learning (ML) and QIS broadly falls into two categories: **1)** classical machine learning for quantum many-body systems and **2)** quantum machine learning (QML). Research in category **1)** has built upon decades of research in ML providing several applications in quantum many-body physics, while category **2)** is a recently developed field, still in its infancy. In category **1)**, we discuss recent progress on the prospects of classical learning with access to quantum data. In category **2)**, we discuss a class of parameterized quantum circuits known as variational quantum algorithms (VQA) which suffer from barren plateaus (BP), making optimization very difficult. In this section, we discuss a new Lie algebraic framework, known as the dynamical Lie algebra (DLA) for characterizing sources of barren plateaus.

Classical machine learning for quantum many-body systems: Over the last half-decade, several classical machine learning methods to learn properties about quantum many-body systems emerged. Given a quantum many-body Hamiltonian, classical machine learning methods can find ground state energies, reconstruct observables, perform quantum tomography, among many other applications. Mathematical results demonstrated that classical machine learning algorithms which learn with access to data in the form of classical shadows, efficiently obtained through few measurements, are proven to be more efficient than those without access to data. In related work, during this program, data from imperfect measurements on neutral atom quantum computers can

enhance variational Monte-Carlo (VMC), demonstrating that hybrid classical-quantum machine learning is approachable within the NISQ era. These results provide new insights into understanding the precise role that data plays in learning quantum many-body systems. Moreover, experimental advances can further enhance classical machine learning methods to explore quantum many-body systems with exotic phases and dynamics.

Variational quantum algorithm methods and training: Variational quantum algorithms are essentially any quantum algorithm that is parameterized and subsequently optimized to minimize some loss function. Given this generic setting, a cornucopia of VQA variations were proposed to suit a variety of different learning tasks.

The most studied form of variational quantum algorithms corresponds to preparing a quantum state that satisfies constraints given by a loss function. Typically this is done by parameterizing a quantum circuit ansatz and applying it to a constant easy-to-prepare state. Sometimes, the goal is to minimize some Hamiltonian that maps to a physical or combinatorial problem, and in other instances it is to find a state that generates a specific distribution when measured (generative modeling). These algorithms all are trained by minimizing a loss function consisting of observables.

There are also QML models that learn a desired unitary, for example . Some of these methods are purely unitary approximations, while others learn a map between input quantum states in a dataset to target output states. Whereas others modify this process slightly by having either the input and/or output data to be classical instead, so they find ways to map the information into/from the quantum system to create an end-to-end process. For instance, one such method presented during the program embeds information about atom configurations into an input quantum state, and then produces a mapping to the associated quantum ground state for that configuration.

Despite the differences between the different types of VQAs, they typically share the same optimization process. A classical optimizer decides how to query the quantum system to find the optimal parameters (though there are ongoing works discussed during this program that are studying how to do this with a quantum optimizer instead). The methods can be 1st or 2nd order gradient based methods, zero-order sampling or global optimization procedures. However, there are a few key differences that make the task more complicated than typical fully-classical optimization. The first issue is that because there is a quantum-to-classical interface, there is essentially always sampling noise. The second issue is that access to gradient information is more expensive than in the case of classical models. On top of this, while the literature in this field is still expanding, it isn't always known how to access analytical gradients for all types of VQAs. So while classical optimization is typically dominated by gradient-based methods, these differences make the choice not so clear-cut for VQAs. As a result, how to best optimize them is still an open question.

Challenges in optimizing variational circuits: A series of works observed exponentially diminishing variance in the loss functions for VQAs of practical use. Referred to as barren

plateaus, this can be understood as the quantum analogue of the zero gradient problem in classical machine learning. Empirical results indicate that as we add more qubits, generic optimization becomes significantly more difficult. A core focus of our working group was understanding the underlying foundations and potential solutions.

The dynamical Lie algebra (DLA) is crucial to our current theoretical understanding of barren plateaus. This is defined as the Lie closure of the generators of the parameterized circuit. The DLA characterizes a sub-algebra of the tangent space of the special unitary group and provides information about the ability of the circuit to produce different operators as the parameters are varied. Recent results presented at IPAM, show that as we increase the dimension of our Lie algebra we reduce the variance of our loss function. This minimization of variance makes optimization difficult, as the loss is concentrated in a small area in the parameter space. Hence with increasing circuit depth and DLA dimension, the variance reduces exponentially and circuits become untrainable. In a less extreme case, low expressibility of the circuit ansatz contributes to barren plateaus because when a circuit has less variation it is difficult to tell whether you are going in the right direction to minimize as you traverse the loss space. Variational quantum methods rely on the gradient or differences between points in the loss space to tell the algorithm whether it is moving towards or away from the minimum.

Work done at IPAM: During the long program, a working group on “Quantum Information and Machine Learning” was formed. This working group organized several research talks from visiting researchers at IPAM, and a reading group where we discussed recent impactful papers. Below is a description of some of the areas we explored.

1. *Empirical Learning Results and Simulability*: One project that came out of these discussions is centered around using empirical results to motivate theory. Specifically, if you can find a unitary to maximize the difference in expectation of an observable applied to a quantum state before and after said unitary is applied, you can use this result to provide bounds on the simulatability of certain groups of operations. While this isn't necessarily quantum machine learning in the sense that we'd expect to do this on a real quantum device, it does involve the learning of unitaries and simulated VQAs to gain insights into where there is a theory-gap in quantum simulation.

2. *Barren Plateaus*: Another interesting direction that was identified at IPAM was to use the tools proposed in the work of unified theory of BPs to help guide the research in the circuit ansatz development for variational quantum algorithms. In the aforementioned work, the authors identified that the building blocks of the Dynamical Lie Algebra, coming from the commutators, can help to determine whether BPs will be encountered in the training. Moreover, they explicitly showed the role that the initial state and the observables play in introducing or mitigating BPs. Interestingly, this could help to identify good circuit ansätze without expensively optimizing the parameterized quantum circuits to understand whether BPs will be encountered.

Outlook: Ultimately quantum learning is an interesting and promising field, but there are many aspects of it that require additional study. On essentially all fronts, there's still work to be done in pushing the limit in both what can be empirically and provably done.

Compared to classical machine learning, quantum models have fewer theoretical guarantees. But especially in the case of VQAs, even if these properties can be shown, there still needs to be research into how we can be confident that we will still be able to properly train them as they scale to larger systems.

7. Open Quantum Systems

Introduction: The mathematical representation for the coupling of a subsystem-bath that accurately represents the response to environmental effects has a recurring presence in many applications in quantum information sciences. Open quantum systems appear naturally in noisy quantum devices, such as photonic and phononic instruments or atom arrays under the influence of an electro-magnetic field. Described in the density matrix picture formalism, the subsystem-environment interactions of interest can be modeled by the Lindblad master equation. The theory of master equations or Lindbladians was originally derived from quantizations of Green functions, and has now emerged as an accepted model to describe the interaction of a small quantum system with a surrounding environment.

Analysis: The time evolution of an open quantum system is determined by a semigroup of quantum channels. In contrast to quantum circuits, these channels require a trace out to the subsystem, and hence produce mixed states. Following the photonic tradition, the generator of such an evolution decomposes in either a dissipative (drive) part and a noisy part given by so-called jump operators, which corresponds to jumps of energy of electrons in an atom array. The analysis of open quantum systems time evolution is of fundamental importance to understand, for example, the convergence speed (mixing time) to the steady state. The Lindbladian evolution fundamentally diverges from the evolution described by the Schrödinger equation. Due to the influence of noise and interactions with the system, the Lindblad evolution is inherently non-reversible and may converge to a steady state at times.

We identified **key problems** for noise models given by Lindbladians: **1)** Determine the steady states; **2)** Determine the mixing time, i.e., the time required to approximately reach the steady state for arbitrary input state; **3)** Describe the characteristic properties (for example periodicity and register swapping) when combining drive and jump part, and how adding drive changes steady states and mixing time; **4)** Identify real experiments for noise simulation; **5)** Determine applications in state preparation; **6)** Find explicit gates for simulation and determine simulation complexity.

Some Lindblad master equations can be equivalently represented by a Wigner-Fokker-Planck formulation, for example in the case of continuous quantum variables. In the Wigner-Fokker-Planck scenario some of the aforementioned problems of interest have been solved. Indeed, for the benchmark problem of the harmonic potential the mixing time is already known by means of studying its Wigner formulation, i.e., it depends on the quadratic form representing the coherent steady state solution and it is obtained by studying the relative entropy of the system in the Wigner formulation. This steady state

solution is known to be independent of the Hamiltonian of the system (as opposed to the solution of the classical Fokker-Planck equation).

More generally, for hermitian jump operators, a whole wealth of techniques is available thanks to tools from representation theory, logarithmic Sobolev inequality and noncommutative transport theory, and recent advances in learning open quantum systems. The following paragraphs will provide a more detailed description of Lindbladian simulation and ground state preparation.

Lindbladian simulation: As the demand to comprehend open quantum systems persists and novel Lindblad-based algorithms emerge, the significance of quantum Lindbladian simulation has grown substantially, gathering considerable attention in recent years. Given a Lindblad equation and a stopping time, a quantum simulation algorithm often refers to a quantum channel with the following two characteristics: 1. With a density operator (or state representation) input, it produces the solution (or state representation) at the stopping time; 2. The circuit can be efficiently implemented on the quantum computer assuming the oracle access to the Hamiltonian and jump operators. Given the non-unitary feature of Lindblad dynamics, the design of this simulation quantum channel deviates from classical Hamiltonian simulation, prompting the need for inventive ideas and techniques in its development.

Ground state preparation: One potential application of quantum computers is the simulation of ground state (GS) properties in quantum many-body systems. However, when assessing the end-to-end algorithmic cost, a significant conceptual bottleneck arises in the form of the state preparation problem. Most quantum systems in nature reach the ground or thermal state exponentially fast. Inspired by this intuition, algorithms that emulate this cooling process might be good candidates for the ground state preparation. Specifically, by constructing jump operators that only allow transitions from higher to lower energies, the system state can be “shoveled” towards the ground state. In recent works, such algorithms were developed with optimized asymptotic complexity and circuit depths, as well as minimal number of ancilla qubits. Although these algorithms behave efficiently in small systems, further numerical experimentation is essential for a comprehensive understanding of their capabilities in larger contexts, see *key problems* 1,2,5.

An important aspect that numerical investigations could assess is the resource cost of GS preparation algorithms associated in terms of circuit depth or number of T-gates with the performance boost achieved in terms of overlap provided from the GS preparation algorithm with the true ground state of the system. Careful resource estimation analysis should be performed to take into account this trade-off and explore the potential speedups associated with ground state energy estimation algorithms given heuristic quantum state preparation algorithms, like VQE, low-depth booster algorithms or Lindblad-based algorithms.

Barren plateaus in the presence of noise: An open question in the field is the connection to open systems that could further be used to characterize the trainability of variational quantum algorithms (VQA) under the effects of noise, see Quantum Learning

section. Thanks to the geometric nature of noise models they are naturally connected to Lie algebra techniques. Through the efforts to characterize the underlying causes of barren plateaus (BP), recent works in Dynamical Lie Algebras (DLAs) provide a mathematical framework to examine the phenomena of BP's in the presence of noise. Indeed, barren plateaus impose a fundamental challenge in QML. A possible connection between the DLAs and open quantum systems could be made by considering the full Hamiltonian structure of the universe, i.e., subsystem and environment, in which the subsystem is embedded in order to have a possibly unitary description of the system including noise.

Work done at IPAM: During the time at IPAM, at least *six projects with different subsets of participants* have been formed to address problems 1-6 mentioned in the Analysis section. There are indeed several papers in preparation concerning simulation time and work on numerical verification. Problem 4 has been addressed in a talk by two experimentalists from UCLA, which will lead to an interdisciplinary work in error correction.

In one of the collaborations a new Lindbladian simulation algorithm was introduced during this program. The new method offers a systematic approach to obtain a unitary approximation to Lindblad dynamics at arbitrarily high orders. Furthermore, the unitary approximation can be efficiently and robustly simulated, involving only Hamiltonian simulation and the tracing out of ancilla qubits, without the necessity for dedicated parameter choices to ensure a high success probability. A closely related project with a different group of collaborators identified concrete gates simulating Lindbladians starting from a fixed set of building blocks. This led to a notion of simulation gap. Using previous work on convex complexity functions, upper and lower bounds it is now in reach to conjecture that simulation cost grows linear in the number of qubits and polynomial in time. In comparing different simulation models, it is paramount to clarify different assumptions for computability.

Complementary to the theoretical endeavors above, numerical simulations emerged as a necessary tool. We developed numerical schemes for simulating the ground state preparation algorithm, which reduced a quadratic time and memory scaling to linear scaling, for implementing the filtered jump operators.

The quantum nature of the Wigner-Fokker-Planck equation prompts work in quantum algorithms for its simulation. A novel avenue is to develop and analyze quantum computing algorithms for this problem and to possibly identify the particular structure of gates to solve the Wigner-Fokker-Planck equation, and compare it to the complexity of the classical counterpart.

In connection to experimental work for prolonging the lifetime of bosonic quantum architectures, photonic error-correcting schemes are being studied for experimental realization exploiting dissipative operators.

Finally, further development of heuristic GS quantum algorithms could play a significant role in improving the performance of many quantum algorithms. Concrete examples of applications of the DLA formalism for practitioners incorporating noise have yet to be identified. Our collaborative efforts with the Quantum Learning group have been instrumental in addressing the challenges posed by this problem.

Outlook: Throughout this program the theory of open quantum systems served as a unifying principle combining analytical, numerical and experimental investigations. The perspective of finding and learning new quantum gates for a myriad of tasks is exciting. This turns a new page in quantum computation. In addition, the broad interest in open quantum systems guides us towards interdisciplinary collaboration on fundamental noise models.

8. Noisy Intermediate-Scale Quantum (NISQ) Computation

Introduction: In the seminal paper from Preskill in 2018, the term *noisy intermediate-scale quantum* (NISQ) computers was introduced. Even though the definition of the term is not very strict, it characterizes quantum devices with a small number of qubits and circuit depth where quantum error correction is too expensive to be implemented on hardware. Since then, the NISQ era has seen significant progress in reducing the error rates characterizing the quantum hardware. Numerous experiments demonstrate that current quantum devices, in spite of all their problems, could still perform certain tasks in this NISQ regime. Of course, the pressing question is whether this era could unfold practical quantum advantage, i.e., showcase that these devices are capable of performing tasks of industrial relevance by either aiding classical computers or directly targeting tasks beyond the capabilities of classical computation. Interestingly, this redirects the research focus back to quantum algorithm developers to come up with problems of practical relevance that current quantum devices could tackle in the range of ~ 100 qubits and ~ 100 circuit depth.

In the following, we further discuss quantum hardware and software developments, *quantum error mitigation* (QEM) techniques, and quantum algorithms suited for NISQ devices.

Hardware and software developments: Multiple quantum computing hardware architectures and paradigms exist. Superconducting architectures implement each qubit as a superconducting circuit. The logical state of a qubit is determined by the number of Cooper pairs which tunnel through an insulating barrier in the circuit, which is manipulated using microwave pulses. Neutral atom and trapped ion architectures use optical tweezers, electric and magnetic fields to manipulate atoms. Photonic quantum computers use photons to represent and process quantum information. Many other architectures exist, each with their own strengths and drawbacks. However, those described above are used

by the largest companies researching quantum computing, including Google, IBM, Microsoft, Xanadu, and IonQ.

Current NISQ-era quantum hardware is restricted by short coherence times, significant gate infidelity, environmental noise, and a limited number of qubits. The first two of these constraints may be addressed by designing shorter and more accurate gates. In superconducting architectures (and others), this can be done via pulse engineering using quantum optimal control. In this approach, an optimization process is performed to find pulses which implement logical gates by minimizing key properties of the gate, such as its infidelity, duration, or the energy of the pulses used in the implementation.

Alongside extensive hardware development efforts, a growing number of open-source quantum computing tools have been proposed in recent years. Developments of the universal protocol framework QASM (now cQASM 1.0) enable defining quantum circuits, gate operations and measurements in the way that can be processed by current quantum computers. To simplify circuits design process, a graphical quantum circuit simulator Quirk is very useful to easily build and manipulate circuits using drag-and-drop features. A variety of software frameworks were created by the major players in quantum computing, e.g. Qiskit (IBM), Cirq (Google Quantum AI), PennyLane (Xanadu), QULACS (Kyoto University, QuanSys) or TKET (Quantinuum), to mention just a few.

Variational Quantum Eigensolver (VQE): The VQE is a hybrid quantum-classical algorithm using the variational principle for ground state preparation of a Hamiltonian, which is one of the most promising algorithms for NISQ devices to address quantum chemistry, materials science, quantum simulations, or optimization problems. A parameterised ansatz circuit is used to produce trial wave functions on the quantum device, allowing efficient measurement of the Hamiltonian expectation value, and classical algorithms are used to variationally optimize the ansatz parameters. Importantly, the VQE as well as other variational quantum algorithms (such as QAOA mentioned in section 5) has been shown to present some degree of resilience to the noise in current quantum hardware. However, the quality of a VQE simulation is only as good as the ansatz itself, and much work has been done to design compact, hardware-efficient and/or adaptive ansätze with the goal of providing high accuracy with few parameters and shallow circuits.

Despite recent advances and successes, a number of bottlenecks have been identified which could prevent the VQE from achieving quantum advantage, such as the substantial cost of observable sampling, exponential scaling due to the barren plateau problem for ansätze with high expressibility, the complexity and convergence of the optimization process, and the need for effective error mitigation techniques needed to improve the accuracy of VQE. These questions are the topic of fast-moving research and have been discussed in the Beyond VQE working group.

Error mitigation: To make near-term quantum algorithms resilient against errors in quantum hardware, various error mitigation techniques have been developed. These methods reduce the error in expectation values without significant increase to circuit

depth, by classically post-processing an increased number of measurements. One method of particular interest to the program was *Probabilistic Error Cancellation*, which probabilistically introduces additional gates based on a learned noise model in such a way that the noise will cancel in expectation. Another error mitigation technique discussed was *Zero-Noise Extrapolation*, where the noise is deliberately amplified with the goal of fitting a noise model to the noisy expectation values and extrapolating to the zero-noise limit. Both these methods rely on accurate models of noise, which can be highly resource-intensive. However, recent advances in modeling noise across the full device has enabled scaling these methods to suppress noise in algorithms on a significant number of qubits. While these were the main error mitigation techniques discussed, other error mitigation methods include filtering results based on known symmetries of the quantum circuit, calibrating qubit readout errors, purification-based techniques, and dynamical decoupling.

Work done at IPAM: During this program, data from experiments performed using a variety of hardware architectures was discussed by speakers from Google (superconducting qubits), IBM (superconducting qubits) and Harvard University (neutral atom qubits). The use of high-order time-stepping methods to compute optimal pulses more efficiently have been investigated within the duration of the program. These could make it practical to directly implement three-qubit gates, or even larger gates, which are more difficult to simulate and consequently more difficult to find optimal pulses for. Three-qubit gates, for example the Toffoli gates, are typically implemented as a circuit of several two-qubit gates. The combined circuit often has a longer duration and higher infidelity than when the pulses to realize the gate are computed directly by pulse engineering. Consequently, this work could help to mitigate the limitations of NISQ era devices. We also had a stimulating group discussion with Prof. Jens Palsberg (Computer Science, UCLA) on current and future efforts, as well as challenges regarding implementations of key quantum algorithms available up to date.

Further, we discussed ways in which variational quantum algorithms can be combined with classical methods for quantum chemistry, and proposed to investigate the effectiveness of these combined methods in the presence of noise from present-day devices. As mentioned earlier, one bottleneck of VQE algorithms is related to the necessary number of measurements. Previous estimates show that for a single energy estimation close to chemical accuracy of relatively small molecules, like methane, the runtime is approximately two days under certain assumptions. Interestingly though, in certain cases, one could use VQE algorithms to get good ground state approximations (with fidelity ~ 0.9) without being within chemical accuracy. In cases where this happens, important runtime reductions for a single energy estimation could be achieved. For example, one hundred times fewer measurements could bring the runtime for a single energy estimation from two days to two hours. Of course, due to the heuristic nature of VQE algorithms and arguments from numerical analysis, there is no guarantee that such a favorable scenario could be realized in all cases, but it is interesting to perform more heuristic simulations and identify use cases where this could be achieved.

Outlook: Whether or not quantum advantage will be possible on near-term (pre-fault tolerant) devices, performing experiments on NISQ devices helps to bridge the gap between hardware capabilities, software developments, and algorithm design. To this end we are interested in building a small toolbox/toolkit that will implement some of the key existing algorithms and testing them on various hardware platforms. Additionally, the high order time-step pulse engineering techniques will be implemented and disseminated as an open source package to be released in 2024.

9. Fault-Tolerant Quantum Computation

Introduction: Many of the quantum algorithms described in this document and throughout the literature require access to stable qubits and high-fidelity quantum operations. Current hardware does not meet these conditions and has error rates many orders of magnitude higher than those of classical hardware. However, in parallel to the development of classical computing, researchers are designing schemes called “quantum codes” that encode logical qubit information across multiple physical, noisy qubits in a way that allows errors to be detected and corrected. Exciting progress is being made in constructing these codes and proving that they have desirable properties, as will be discussed further in this section. Quantum error correction (QEC) also has interesting connections to other areas of computer science, mathematics, and physics, some of which are briefly mentioned. As quantum computing progresses towards fault tolerance, most expect that we will pass through an era called “early fault-tolerant quantum computing” (EFTQC) which has unique constraints for algorithm design and has garnered recent attention. It will also be important for progress towards fault tolerance to estimate the end-to-end resources required to run algorithms in this regime.

Quantum error-correcting codes: There are many categories of quantum error-correcting codes, including but not limited to surface codes, low-density parity check (LDPC) codes, subsystem codes, and Floquet codes. These can be formulated in terms of stabilizers, where each code has a group of mutually commuting stabilizers that leave the quantum computer’s initial state unchanged. Any errors that do not commute with the stabilizers are then detected upon measurement of a corresponding ancilla qubit.

Kitaev’s toric code uses parity check ancilla qubits to measure local X or Z stabilizers on four neighboring qubits along with a periodic boundary condition (as in the torus). Implementing the toric code with the periodic boundary condition may instead involve non-local operations, rendering it unfavorable for current devices. On the other hand, the surface code has identical formalisms but without the periodic boundary condition. While both of these topological codes can achieve an arbitrary code distance (scaling as square-root of the number of physical qubits), the number of encoded logical qubits remains constant, resulting in a large physical qubit overhead for more logical qubits.

A long-standing open question has been to find quantum codes with good properties, such as ensuring that both the number of encoded qubits and the code distance scale linearly in the number of physical qubits. Recently, this feat was achieved through the

development of quantum LDPC codes, derived from expander graphs. These codes also have connections to the quantum probabilistically checkable proofs (PCP) conjecture, which is the quantum counterpart to the classical PCP theorem. Unlike the PCP theorem, quantum PCP is still a conjecture. However, a recent resolution to the no low-energy trivial states (NLETS) conjecture was found using these quantum LDPC codes, implying there could be hope to resolve quantum PCP affirmatively in the future.

Early fault-tolerant quantum computers: Overheads for fault tolerance seem to present a particular challenge as we transition from the NISQ era to the era of early fault tolerance. Algorithms research in other sections of this white paper has focused on reducing the number of ancilla qubits required to implement certain subroutines (e.g., energy estimation or state preparation). Nevertheless, other considerations will likely augment these constraints as hardware continues to mature. For example, demonstrations on early fault-tolerant hardware have already primarily focused on implementing non-universal transversal gates. In fact, by the Eastin-Knill theorem it is clear that transversal gate sets *cannot* be universal and additional logical resources are required to implement universal quantum computation, e.g., magic state distillation. Thus the realization of the first universal fault-tolerant quantum computer will likely be dominated by the cost of methods like magic state distillation. The capabilities of early fault-tolerant quantum computers and the prospects for using them to achieve quantum advantage are challenged by these overheads, and the notion that they will likely have vastly different logical error rates for Clifford and non-Clifford operations. Bespoke algorithms that optimally account for these asymmetries are likely to be an area of active research in coming years.

The early fault-tolerant quantum computing (EFTQC) regime could help better understand how the transition from NISQ to FTQC could be realized in practice. For example, studying the interplay between QEM and QEC can characterize the transition from NISQ-to-EFTQC. Introducing realistic noise models that capture the performance of quantum devices can guide algorithmic developments to exploit the capabilities of quantum hardware. Another interesting open direction is the transition from EFTQC-to-FTQC. To this end, benchmarking quantum algorithms that traditionally belong to EFTQC and FTQC could help understand what problem sizes in terms of number of logical qubits each algorithm could tackle given the same number of physical qubits and error correction scheme. The EFTQC era could help bridge the gap between NISQ and FTQC regimes and the corresponding efforts in these two fields of quantum computation.

Fault-tolerant resource estimation: The overheads associated with enabling fault-tolerant quantum computation are among the most significant challenges confronting the realization of quantum computers capable of executing algorithms that might achieve quantum advantage. The large number of physical qubits needed to encode logical qubits is the overhead that typically receives the most focus, but the implementation of logical operations, decoding, routing, and magic state distillation comprise others that contribute to the hardware requirements for implementing any given algorithm. There is a growing body of research focused on quantifying these costs for different codes and architectures, with the aim of understanding which quantum algorithms (and input instances) can be expected to be implementable in computational volumes that are motivated by optimistic

assumptions about what advances in physical hardware might be realized in the coming years.

Presently, the majority of analyses have focused on relatively conventional surface codes and the assumption that physical error rates will eventually reach one error per thousand or ten thousand physical operations. In that context, specific challenge problems in quantum chemistry (e.g., sampling from the energy eigenspectrum of FeMoco) have converged on resource estimates in the neighborhood of thousands of logical qubits with logical error rates sufficient to execute tens of billions of non-Clifford operations (i.e., the dominant cost associated with fault-tolerant quantum computation in these architectures). These resource requirements are presently vastly larger than any available NISQ hardware and bringing these costs down through innovation in algorithms and architectures are great motivators. It also warrants noting that there are relatively few detailed end-to-end analyses of these costs, though we anticipate that this will change rapidly as progress in implementing fault-tolerant quantum computation is made.

ZX-Calculus: ZX-calculus is a graphical language used to describe linear maps between qubits. It emerged from categorical quantum mechanics and consists of string diagrams paired with rewrite rules that allow computations to be done graphically. Under certain rewrite rules the ZX-calculus has been shown to be sound, universal, and complete. These properties serve as the foundation for a growing area of research that uses ZX-calculus as a tool to analyze and better understand certain aspects of quantum computation. Such works have had significant impacts in areas including circuit compilation and optimization, ansatz interpretation, measurement-based quantum computation, and quantum error correction (QEC).

The influence of ZX-calculus is extensively growing in the QEC literature for several reasons. ZX-calculus is especially well-suited for tasks that require an understanding of gate-level quantum operations, which is currently the way quantum codes and checks are described. ZX-calculus also describes general linear maps between qubits, rather than just unitary maps. Thus, it can be used in instances where representations using typical quantum circuit diagrams might be inadequate. For example, it is used extensively to design, verify, and optimize lattice surgery procedures for the surface code, since basic lattice surgery operations act non-unitarily on logical states. Another success of ZX-calculus in QEC was the recent unification of different fault-tolerance paradigms. The authors showed that circuit-based, measurement-based, fusion-based, and Floquet-based quantum computation all share an underlying stabilizer fault-tolerance structure using ZX-calculus. These results give researchers a way to “translate” between the different paradigms and hopefully enable better transfer of ideas and successes between the different areas.

K-theory, geometry and codes: Quantum mechanics is based on Hilbert space and the linear algebra thereof. Freedman calls this a naked Hilbert space. Quantum Information Science requires additional structure, in particular tensor product structure. Quantum entanglement highlights how such additional structure brings in surprising implications and fertile applications. On the other hand, Hilbert spaces can be endowed with manifold-

like structures. Synthesizing both tensor and parameter structures to the naked Hilbert space has yet to be fully investigated. Recent work by Freedman and Hastings suggests that K-theory of quantum information or quantum K-theory is the emerging structure. This new “quantum K-theory” should inherit the regularity of traditional K-theory such as Bott periodicity. This would indicate hidden regularity structure for entanglement, which would come to life through codes and error correction.

Recent breakthrough construction for quantum LDPC codes stems from a product of classical LDPC structure and combinatorial data of expander graphs. Vice versa, according to Freedman and Hastings, codes can be used to construct interesting eleven-dimensional manifolds. This corroborates the idea that homological and geometric properties of manifolds could productively inspire construction of codes with either better decoding properties or additional more exotic features. This also applies to the overlap with black hole physics, where information of a black hole is encoded into the boundary of a quantum field theory. The information collected during the long program suggests that a deeper analysis is in order.

Work done at IPAM: During this program a group formed to work on constructing explicit examples of small quantum Tanner codes. Although quantum Tanner codes are a provably good LDPC construction in the asymptotic limit, there are no known explicit constructions with good properties. We reviewed the code construction and selected the smallest field over which to construct left-right Cayley complexes, resulting in a code requiring about 700 physical qubits. Work is ongoing to implement this example and determine the number of logical qubits encoded, as well as to empirically estimate the code distance.

Another project completed during the program was an analysis of the fault-tolerance overheads associated with implementing different energy estimation algorithms intended for early fault-tolerant quantum computers. It was found that phase estimation algorithms based on the Hadamard test generally have lower overheads than iterative phase estimation, thanks largely to the relative robustness of these algorithms to depolarizing logical errors.

An additional group formed to study ZX-calculus and some of its applications. Although not so many of the QEC connections were explored, the group spent significant time working with an extension of ZX-calculus, called ZXW-calculus, which makes use of W states to facilitate exponentiation of Hamiltonians. Using ZXW-calculus, the group obtained an exact Magnus remainder of the product formula for unitary propagators of several Hamiltonians.

Outlook: Fault-tolerant quantum computers are necessary for performing truly useful calculations to impact fields from cryptography to chemistry. Although progress in quantum error correction is steady, more development is needed to realize fault-tolerant quantum computation. More ideas and tools are utilized and integrated from diverse fields, ranging from AdS/CFT holography, sub-systems, to non-invertible symmetries. Tensor categories will continue synthesizing diverse approaches to QEC. Quantum K-

theory needs to be further understood, but it is already clear that it sheds new light on the way we address QEC. While the ZX-calculus and its variants are emerging as tools providing a different perspective on known results, we anticipate broader adoption of and familiarity with the ZX-calculus will enable new kinds of results.

10. List of Acronyms

AdS/CFT	Anti-de Sitter/conformal field theory
BE	Block encoding
BP	Barren plateau
BQP	Bounded-error quantum polynomial time
CSP	Constraint satisfaction problem
DLA	Dynamic Lie algebras
EFTQC	Early fault-tolerant quantum computation
FTQC	Fault-tolerant quantum computation
GS	Ground state
LCU	Linear combination of unitaries
LCHS	Linear combination of Hamiltonian simulation
LDPC	Low-density parity check
ML	Machine learning
NISQ	Noisy intermediate-scale quantum
NLTS	No low-energy trivial states
ODE	Ordinary differential equation
OQS	Open quantum system
PCP	Probabilistically checkable proofs
PDE	Partial differential equation
QAOA	Quantum approximate optimization algorithm
QEC	Quantum error correction
QEM	Quantum error mitigation
QFT	Quantum field theory
QHD	Quantum Hamiltonian Descent

QIS	Quantum information science
QLSA	Quantum linear system algorithm
QMA	Quantum Merlin Arthur
QML	Quantum machine learning
QRAM	Quantum random access memory
QRACM	Quantum random access classical memory
QRAQM	Quantum random access quantum memory
QSVT	Quantum singular value transformation
SAT	Boolean satisfiability problem
SDP	Semidefinite programming
VQA	Variational quantum algorithm
VQE	Variational quantum eigensolver