

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.



Quantum Information for Fusion Energy Sciences

Final Technical Report

Rigetti Computing,
Lawrence Livermore National Laboratory,
University of Southern California

DOE-RIGETTI-0021661

<http://www.rigetti.com>

22 December 2024

Preamble

Purpose

This document serves as the Final Scientific and Technical Report in accordance with Assistance Agreement Award No. DE-SC0021661, Federal Assistance Reporting Checklist (F). Part of this work (LLNL-TR-2000784) was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Acknowledgement

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Fusion Energy Sciences, under Award Numbers DE-SC0021661, SCW-1680 and SCW-1736. Work by Lawrence Livermore National Laboratory was performed under the auspices of the U.S. DOE under contract DE-AC52-07NA27344.

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Executive Summary

The simulation of plasma dynamics is a critical area of Fusion Energy Sciences (FES) due to its usefulness in predicting, controlling, and confining plasmas in the context of potential fusion reactors. The simulation of plasmas is a computationally difficult problem in both classical and quantum physics, motivating investigation into the potential of quantum computers to simulate these systems. This project took several concrete steps towards this goal by developing tools for improving the control, characterization, and calibration of quantum gates on a superconducting quantum computer, developing error suppression and mitigation tools to reduce errors on the quantum computer, and utilizing these advancements to simulate reduced models of plasma dynamics on the quantum computer.

In order to efficiently simulate plasma physics, an optimal control method which synthesizes, directly at the pulse level, any quantum gate on qubit and qutrit systems was developed. Using four superconducting transmon quantum processors at Rigetti and LLNL, it was demonstrated that any arbitrary quantum gate on qubits and qutrits could be implemented with high fidelity, leading to a significantly reduced length of a gate sequence [1].

A problem of interest in FES is the nonlinear optical process of laser pulse compression within a plasma. Since quantum physics is linear, simulating nonlinear operations is not naturally feasible on a quantum computer, however it is possible to simulate a quantized version of the nonlinear process. A quantization approach to convert nonlinear wave-wave interaction problems to Hamiltonian simulation problems was developed and demonstrated using two qubits on a Rigetti device. In this experiment, a number of error suppression and mitigation techniques were investigated to determine how best to utilize the finite quantum resources. This study provides an example of how plasma problems may be solved on near-term, noisy quantum computing platforms and identified a promising set of techniques [2].

Building on the insights of these experiments, the investigation turned to linear electron-plasma wave physics. A connection was identified between a local one-dimensional lattice spin model and linear wave phenomena, allowing a plasma physics problem to be efficiently mapped to the quantum computer. In this framework, reflection and transmission of plasma waves at a sharp boundary was studied, as well as the propagation of waves through an inhomogeneous plasma medium. In addition to the suite of error suppression and mitigation techniques developed, this experiment introduced the use of a digital-analog gate scheme designed to efficiently simulate the plasma Hamiltonian. With hardware available at the conclusion of the project, simulation at the scale of 9 qubits and 15 timesteps (60 entangling layers) was achieved.

Contents

1	Introduction	1
1.1	Abstract	1
1.2	Overview	2
1.3	Background	3
1.3.1	Towards Quantum Advantage for Plasma Physics Applications	4
1.3.2	Operating Superconducting Qubits at Noisy Intermediate Scales	6
1.4	Results	8
1.4.1	Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits	8
1.4.2	Simulating nonlinear optical processes on a superconducting quantum device	9
1.4.3	Simulating electron-plasma waves on a superconducting quantum chip	10
1.5	Conclusion and outlook	11
1.6	Acknowledgments	12
2	Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits	13
2.1	Overview	13
2.2	Method	13
2.3	Discussion	18
2.4	Conclusion	20
3	Simulating nonlinear optical processes on a superconducting quantum device	21
3.1	Classical model of laser pulse compression	21
3.2	Quantum model of laser pulse compression	22
3.3	Exact quantum dynamics	24
3.4	Implementing exact dynamics with error mitigation	25
3.5	Testing product formulas and optimal use of quantum resources	28
4	Simulating electron-plasma waves on a superconducting quantum chip	32
4.1	Introduction	32
4.2	Theory	33
4.2.1	Electron Plasma Wave Model	33
4.2.2	Ion Sound Wave Model	35
4.2.3	Spin model	37
4.3	Electron-plasma wave simulations with the spin model	37
4.3.1	Experimental scheme to measure the spectrum	37
4.3.2	Experimental scheme to simulate injection of plasma waves	38
4.3.3	Propagation in an inhomogeneous medium	39
4.4	Running an experiment on Ankaa-9Q-3	40

Contents

4.5 Experimental results	41
4.5.1 Dispersion	42
4.5.2 Scattering of Waves in an Inhomogeneous Plasma Medium	43
4.5.3 Reflection of Waves from the Plasma Edge	44
4.6 Conclusion	45
Appendix A Supplemental Material	46
A.1 The PMW3 decomposition	46
A.2 Error budget	47
A.2.1 Device design and architecture	48
References	50

1

Introduction

1.1 Abstract

The overarching goal of this project is to simulate reduced models of plasma dynamics using advanced control techniques for the exponential Hilbert space of intermediate-scale quantum information processing (QIP). Discovering methods for the quantum acceleration of plasma simulation techniques is important for Fusion Energy Sciences (FES) because predicting, controlling, and confining the fourth state of matter includes a number of challenging problems in both classical physics and quantum physics. Hence, quantum simulation of plasma dynamics, as well as nonlinear system dynamics and field theories in general, has become an important target for the realization of quantum advantage with QIP.

This project investigated the possibility of achieving quantum advantage through the co-design of fusion plasma applications and optimal control architectures for advanced experimental hardware systems based on multi-qubit superconducting devices. By conducting research within three key thrust areas of theory, control, and experiment, we investigated how to efficiently map plasma models onto quantum memory registers and how to emulate the cubic and quartic nonlinear interaction terms that govern the physics of plasmas by using coherent qubit-qubit entangling interactions in superconducting circuits. These interactions are native to devices but otherwise unexploited for traditional QIP. This project explored the development and application of control pulse engineering techniques to the Rigetti platform in order to enable long time simulations with high effective gate depth. High fidelity was achieved by developing and applying dynamical error suppression techniques to the quantum simulation of plasma dynamics.

The outcome of this project was the first exploration of carefully engineered multi-qubit gates and multi-qubit interactions for the simulation of plasma dynamics. The combination of advanced control and error suppression techniques allowed implementation on Rigetti platforms composed of 9 qubits, representing Hilbert spaces of 512 dimensions. Hence, this project accelerated the identification and development of quantum simulation methods for fusion plasmas that have the potential to exceed classical computing capabilities and laid the groundwork for quantum advantage over a 5-10 year horizon. It also led to a significant expansion of control libraries for the open-source QUIL framework that will benefit other research domain areas beyond FES. Thus, our collaborative project acted as a bridge between the QIS and FES communities that will enrich the scientific ecosystem.

1. Introduction

1.2 Overview

Quantum information science (QIS) provides a number of potentially transformative new approaches for understanding, designing, and ultimately achieving, controlled fusion energy systems [3]. In particular, plasma simulation [4] is especially important for advancing towards next generation fusion systems by enabling system designers and operators to understand and predict behavior with high-fidelity computational models [5]. Hence, as discussed in Priority Research Opportunity (PRO) 1 of the 2018 FES Roundtable Report on QIS [3], discovering methods for quantum acceleration of plasma simulation techniques is important for fusion energy sciences (FES) since diagnosing, controlling, and confining the fourth state of matter includes challenging problems in both classical physics and quantum physics. Thus, quantum simulation of plasma dynamics, [6, 7, 8, 9], and of classical dynamical systems and field theories [8, 9], represents an important target for the realization of quantum advantage, and thus quantum simulation of reduced models of plasma dynamics [6, 10] have emerged as an important research focus area in quantum computing. Moreover, as explained in PRO 6 [3], advanced theoretical and computational capabilities within the FES community, especially for the control of complex nonlinear dynamical systems, are likely to contribute valuable new approaches to the optimal control of quantum information processing (QIP) systems which are also complex nonlinear systems that are dissipatively coupled to their environment.

This project investigated how to perform Hamiltonian simulation of reduced models of plasma dynamics in an efficient manner using advanced quantum computing hardware platforms and advanced methods for co-design of quantum hardware and software systems. We addressed PRO 2 by exploring how to efficiently map plasma models to quantum memory registers and how to emulate the cubic and quartic nonlinear interaction terms that govern the physics of plasmas by using coherent qubit-qubit entangling interactions in superconducting circuits (see Fig. 1.1). Optimal control techniques were used to develop quantum logic that can efficiently simulate nonlinear Hamiltonian interactions with high effective gate depth and, hence, high effective fidelity relative to standard approaches to quantum gate sets [11, 10]. Open-loop dynamical error correction methods, including dynamical decoupling (DD) and dynamically corrected gates (DCG), were used to suppress and correct errors in a generic and modular fashion that can be used to further improve fidelity and can be applied both when QPUs are otherwise idle (DD) or active (DCG). Our objective was to simulate models of plasma dynamics on QIP lattices composed of many qubits. In achieving this, we significantly accelerated the identification and development of quantum simulation methods for fusion plasmas that have the potential to exceed classical computing capabilities and demonstrate quantum advantage over the next 5-10 years.

Our project was informed by early lessons learned from the LLNL-Rigetti collaboration [10] which demonstrated relative advantage for LLNL's optimal control-based approaches [11] over traditional gate-model QIP architectures for near-term plasma applications on systems with only a few qubits. Upon moving to an intermediate scale

1. Introduction

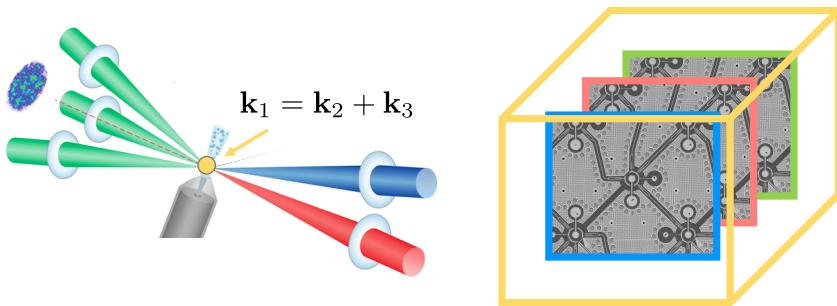


Figure 1.1 – Recreating Plasma Dynamics on an Intermediate Scale Quantum Processor: [Left] Parametric interactions are key ingredients for fusion. Even with low-orders of mixing terms, these effects can give rise to complex dynamics, as observed for instance, in cross-beam laser energy transfer. There, ion acoustic waves (indicated in gold) mediate parametric three-wave mixing between high-intensity laser fields. [Right] This work considers the simulation of these effects by controlling targeted quantum-coherent processes within the exponential Hilbert space of a superconducting multi-qubit processor. Figure adapted from [Tunbull2020, 12].

on experimental Rigetti devices, we anticipate that these advantages will compound, making plasma physics an important test case for designing near-term QIS systems. As described in Table 1.1, our collaboration brought together experts spanning fusion plasma theory, quantum control and architecture co-design, quantum error correction and computation, and superconducting device theory and NISQ integration.

Institution	Key Personnel	Key Expertise and Research Areas
Rigetti	B. Evert (PI) E. Sete (co-PI) X. Wu	Quantum computing systems with superconducting qubits Quantum device theory Hardware architecture co-design for fusion applications
LLNL	Y. Cho (co-PI) V. I. Geyko I. Joseph	Quantum information theory and quantum optimal control Plasma physics Fusion energy science
USC	D. Lidar (co-PI) Amy Brown	Digital and analog quantum computation Error correction and mitigation

Table 1.1 – Collaboration: Our project leveraged a multi-sector, interdisciplinary team with expertise in the codesign of superconducting quantum hardware for plasma physics applications.

1.3 Background

The computational challenge towards predicting plasma dynamics with high accuracy is immense. One potential path forward is with quantum simulation [13]. Emerging technologies for quantum simulators range from specialized experimental setups to fully

1. Introduction

universal, gate-model machines. Universal processors from Rigetti, based on superconducting qubits, have enabled domain-experts to explore early, digital QIP applications spanning nuclear, condensed-matter, chemistry, and material sciences, and notably, the first experimental demonstrations of QIP for plasma simulation [10], which were conducted by members of our team. In this project, we considered hardware-tailored QIP implementations of nonlinear classical, semiclassical, and quantum plasma dynamics. As a shared commonality among these, the governing equations have nonlinear interaction terms that couple the dynamics of multiple waves and/or particles [14]. Many quantum systems, including superconducting quantum circuits, can also exhibit these types of interactions in coherent dynamics. By studying whether quantum computers can be used to solve three-wave mixing problems by emulating the multi-body dynamics directly in an exponential Hilbert space, we addressed a fundamental class of problems with applications to many nonlinear systems.

Our project parallels a fundamental challenge for fusion systems: controlling extraordinarily complex physical systems with highly nonlinear interactions, involving an enormous expanse of space and time scales. For FES, this has already led to the development of simulation tools to support some QIS systems, such as ion traps [3]. We anticipate that the computational tools already used by FES domain experts or co-developed based on expertise in the FES community could prove useful for QIS, and we intend to engage this body of work directly. Specifically, optimal control techniques applied to open quantum systems are expected to be a powerful toolbox, yielding pulse-level improvements to target controlled specific unitary evolution, tailored to theoretical problems in FES. Control is therefore the central pillar of this project.

1.3.1 Towards Quantum Advantage for Plasma Physics Applications

Plasma is the phase of matter in which electrons and ions have enough energy to become unbound from neutral atoms. For example, in many regimes of interest the plasma equation of state is well approximated by the ideal gas law, but with an increase in pressure due to the number of free electrons. Technically, in order to satisfy the definition of a plasma, the density must be high enough that the so-called plasma parameter, the number of particles in a Debye sphere, is much larger than one. This ensures that plasma behaves as a collective system that tends to screen out electric fields on scales larger than the Debye length, rather than as a system of independent particles. Due to the high temperatures and number densities required for optimal fusion power output, ~ 10 keV for D+T fusion reactions, the easiest fusion fuel to ignite, classical models of the electromagnetic fields and particle distribution functions are usually employed. Due to the even larger computational complexity of treating these systems via their fundamental quantum description, quantum plasma models are rarely studied. However, the quantum dynamics of plasma are of strong interest for a number of emerging applications in plasmas under extreme conditions, such as high energy density plasmas, strongly correlated plasmas, and intrinsically quantum plasmas (as described in DPS-D, DPS-E, and DPS-H of the recent community report [15]).



1. Introduction

Even classical plasma physics models, where the electromagnetic fields are taken to satisfy Maxwell's equations, have a high degree of complexity. In the limit of weak interactions, the single-particle kinetic probability distribution function (PDF), a function of six dimensions: three real space and three velocity space coordinates, in addition to time, is usually taken as the fundamental description. In this limit, the PDF satisfies the Liouville equation, also called the Vlasov equation, and interparticle correlations are sufficiently weak that the collision operator, which represents nonlinear interactions between particles, can be determined from an approximation to the two-particle correlation function [16]. In truth, a more fundamental description would include the PDF for all N particles, which would require solving a $6N$ -dimensional partial differential equations. Since solving for the evolution the six-dimensional single PDF function already taxes the limits of the world's most advanced supercomputing facilities, solving such high-dimensional models is clearly unthinkable with present-day capabilities. Yet, simulating and solving high-dimensional partial differential equations are one of the application areas in which future fault-tolerant quantum computers may offer significant advantage [8, 9].

Thus, understanding how quantum computers can be used to simulate the complex nonlinear dynamics of plasmas is of great interest to FES. It has been predicted that quantum computers can achieve exponential speedup over classical computers for the simulation of intrinsically quantum systems [17, 18]. Quantum advantage, typically a quadratic or general polynomial speedup, has also been predicted for the simulation of intrinsically classical systems, including the simulation of linear ordinary and partial differential equations (PDEs). [19, 20, 21], optimization problems [22, 23], and Monte Carlo methods [24]. Thus, research into the possibility of achieving quantum advantage for intrinsically classical computations has been active since the inception of research into quantum algorithms [25].

Recent research has focused on the possibility of achieving quantum advantage for plasma physics applications [6, 8, 7, 9]. Large computational savings are anticipated for high-dimensional PDEs, relative to an Eulerian discretization of the numerical problem, because the exponential scaling of the computational complexity with respect to the dimension of the PDE can be reduced to polynomial scaling. Such techniques may eventually be applicable to high-fidelity kinetic simulation of the six-dimensional (6D) single particle PDF of plasma physics [3]. Moreover, such algorithms could eventually enable the rapid solution of even higher dimensional PDEs, such as that obeyed by the twelve-dimensional (12D) two-particle correlation function (or higher multi-particle correlation functions), which are necessary to accurately evaluate two-particle (and multi-particle) collision operators and turbulent correlation functions in plasmas [8].

There are a number of promising approaches to the Hamiltonian simulation of nonlinear classical dynamical systems, even though idealized quantum computers can only apply linear unitary operations. A natural approach is based on simulating the quantized version of the classical Hamiltonian system, which allows the study of both the quantum and the classical dynamics. A good example of this technique is represented by a new hardware-efficient algorithm for the quantum simulation of nonlinear 3-wave interac-

1. Introduction

tions (see Fig. 1.2), recently developed by the LLNL team as a paradigm for nonlinear interactions in plasma physics, fluid dynamics, and gauge theory [10]. The nonlinear 3-wave mixing problem denotes a generic vertex interaction in which a wave (or particle) of type 1 decays into waves of type 2 and type 3. In order for the three-wave problem to be Hamiltonian, the inverse process of waves of type 2 and type 3 recombining into waves of type 1 must occur with equal magnitude. On a fundamental level, three-field vertices are required to couple fermions with bosons, and mediate boson self interactions (along with four-boson terms) in nonabelian gauge theories. On an applied level, three-wave couplings are responsible for parametric interactions in crystals, turbulence cascade in fluids, and cross-beam laser energy transfer in plasmas (illustrated schematically in Fig. 1.1).

The new quantum algorithm for the three-wave problem uses the integrability of the quantized 3-wave problem to exactly solve for the unitary evolution operator. For this problem, there are two conserved adiabatic action invariants that commute with the Hamiltonian, which implies that the quantized system is integrable. Once the states have been transformed to the action basis, the Hamiltonian becomes a tridiagonal banded matrix and sparse Hamiltonian problems can be simulated efficiently using quantum computing hardware.

However, many fundamental models of plasma physics include the effects of dissipative processes mediated by particle collisions. How can these systems be treated when the underlying dynamics are not Hamiltonian? A general framework has been developed for the quantum simulation of nonlinear classical dynamics on phase space [8]. The Koopman-von Neumann formulation of classical mechanics on Hilbert space has been generalized to treat arbitrary non-Hamiltonian nonlinear classical dynamics. The Koopman-von Neumann representation implies that conservation of the PDF on phase space, as expressed by the Liouville equation, can be recast as an equivalent Schrödinger equation on Hilbert space with a Hermitian Hamiltonian operator and a unitary evolution operator. Hence, a quantum computer with finite resources can be used to simulate a finite-dimensional approximation of the operator corresponding to the Koopman-von Neumann Hamiltonian. Quantum simulation of classical dynamics can be proven to be exponentially more efficient than a classical Eulerian discretization of the Liouville equation if the Koopman-von Neumann Hamiltonian is sparse [8]. Utilizing quantum walks and associated techniques can potentially lead to a quadratic improvement over classical time-dependent Monte Carlo methods [24].

1.3.2 Operating Superconducting Qubits at Noisy Intermediate Scales

The computational power of a quantum algorithm is generally informed both from its width, the number of qubits available for computation, and from its depth, the number of coherent operations that are executed (before the accumulation of error rates overwhelm the computation). For near-term quantum computers, before fault-tolerant machinery, circuit depth is often the limiting factor for a quantum algorithm. Within



1. Introduction

the set of operations needed for universal quantum computation, typically the error rates of the two-qubit gates are the limiting factor in determining the performance of a NISQ-era quantum computer [26]. Gates themselves are an abstraction for calibrated, analog quantum processes.

There are numerous techniques for implementing two-qubit gates in superconducting qubits [27]. Rigetti pioneered the application of parametric processes for two-qubit gates to mitigate cross-talk issues at intermediate scales [28, 29]. For this gate scheme, the resonant frequency of a single transmon qubit is modulated using local flux control. The modulation frequency can be chosen to satisfy energy conservation for a number of processes in a multi-qubit circuit with static capacitive coupling [30]. The resulting process is a frequency-selective exchange interaction, allowing for entangling gates between otherwise weakly-interacting qubits. This gate scheme was implemented in Rigetti Aspen series processors, which were the testbed of the initial experiments.

The fixed-capacitance between the qubits used on the Aspen processors resulted in unwanted residual ZZ interactions during idling periods when single-qubit gates are applied. To address this challenge, tunable couplers were introduced as an effective solution for minimizing such errors while enabling fast two-qubit gates with improved fidelities. Rigetti developed a floating tunable coupler [31] that not only suppresses residual ZZ interactions but also facilitates high-speed entangling gates without relying on direct qubit-qubit capacitance. This innovation has enabled coupling between qubits located on different chips, a key milestone in the realization of modular superconducting quantum architectures. Within this framework, Rigetti has implemented two types of gates: parametric resonance gates and base-band fast DC gates. Additionally, the architecture allows for the activation of alternative interaction mechanisms, such as cross-resonance gates or coupler-modulated gates, further expanding its versatility. This technology forms the basis of the Ankaa-series processors which are used in the final experiments.

Quantum optimal control is one of the leading toolboxes for constructing targeted unitary gate operations [32], which can evolve each time step of a physical simulation of complex dynamics [33]. The first-ever successful simulation of multiple cycles of complex nonlinear plasma dynamics (see in Fig. 1.2) was only achieved on the LLNL Quantum Design and Integration Testbed (QuDIT) quantum computing platform through the use of state-of-the-art techniques for applying optimal control of quantum hardware to the simulation of Hamiltonian evolution [10], leveraging the Hilbert space of a single transmon. In contrast, as seen in Fig. 1.2, simulation of the same problem on two qubits of the Rigetti Aspen-4 quantum processing unit (QPU) was only able to achieve enough depth to simulate the first half period of the dynamics. The intrinsic hardware noise that determines decoherence on both computing platforms is similar, so the majority of the difference is due to the fact that the sequence of standard gates required for each time step uses about 17-18 times as many gates and about 20 times as much time duration. Hence, it is the substantial reduction in time duration afforded by optimal control techniques that is responsible for improving the performance with respect to decoherence. As this case study shows, a co-design approach to optimal control of quantum hard-



1. Introduction

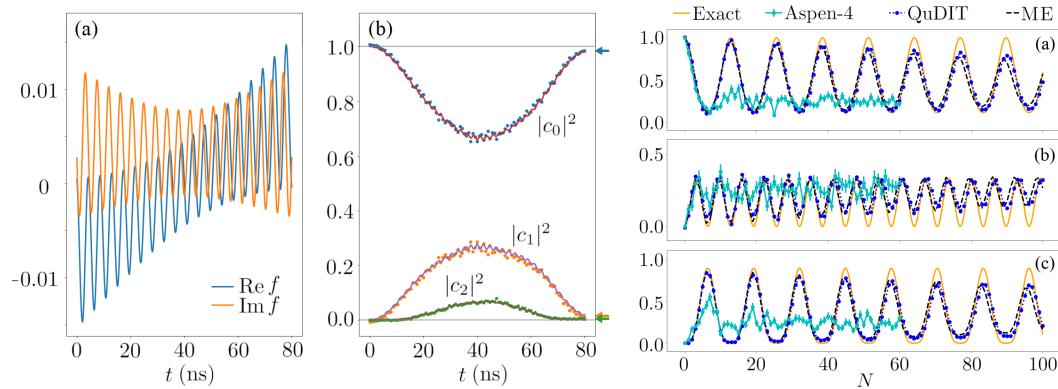


Figure 1.2 – Three-wave interactions with engineered cubic couplings: [left] (a) An optimized microwave control pulse, $f(t)$ was used to recreate a single time step for the three-wave problem using Hamiltonian simulation. (b) Experimental results for the population of each transmon level, $|c_j|^2$ versus time, compared to Lindblad master equation model for the LLNL QuDIT platform. [right] The probability, $|c_j|^2$, of the occupation number $j = 0, 1, 2$ for one wave in the problem versus time step, N . Experimental results for the LLNL QuDIT platform (dark blue) and the Rigetti Aspen-4 chip (cyan), compared to analytics (orange) and Lindblad master equation model (ME, black dashed). Figures from [10].

ware is a highly efficient way to steer a quantum system through Hamiltonian simulation. However, scaling these techniques to many-body interactions across superconducting qubits represents a formidable scientific and engineering challenge.

1.4 Results

The prior work provided a clear research agenda to investigate the potential of near-term quantum computers to simulate plasma dynamics, leading to the development of a set of theoretical and practical techniques which culminated in experimental demonstrations. These demonstrations form the structure of this report, with three experiments performed, each composing a chapter. While each chapter can be read individually, the sequence of the three illustrates the evolving theoretical understanding as well as the developing set of tools for plasma simulation on quantum computers.

1.4.1 Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

Contemporary quantum computers typically provide a universal gateset, where the quantum computer is calibrated with a set of discrete pre-determined gates. The gateset is determined by a balance of what the platform can easily achieve, and what is generally useful for algorithms. However, the universal gateset is unlikely to be optimal for any

1. Introduction

particular algorithm, leading to the opportunity of co-design. What the native gateset could be specially tailored to the problem at hand?

This approach to simulation is often called optimal control. These methods begin with an accurate physical model of the device, and determine a set of control signals that will implement a target unitary. Optimal control promise time- and resource-efficient compact gate sequences and, therefore, higher fidelity.

Moreover, while contemporary devices are focused on qubits with two possible states, higher-dimension “qudits” can offer exponentially greater power. Thus, a set of qutrits, systems with three computational states, were investigated.

This line of work resulted the demonstration of custom pulses capable of generating any target unitary with high fidelity and on a wide variety of devices and qubits. Using both qubits on large-scale Rigetti QPUs and qutrits in the LLNL lab, optimal control gates were demonstrated and benchmarked. The calibration overhead for such gates is manageable, setting the stage for optimal-control based circuits in future work.

The results of this study were published in Physical Review Applied [1].

1.4.2 Simulating nonlinear optical processes on a superconducting quantum device

Building on the work described in [10], investigation into plasma simulation using a standard digital set of gates was continued. Previous results showed that while the problem could be effectively posed as a quantum circuit, the high depth of the circuit meant that errors quickly suppressed the signal of interest.

A set of error suppression techniques were explored, including improving fundamental gate fidelities, improving compilation efficiency, and dynamical decoupling. These resulted in significantly improved error rates for the circuit. However, even with lower error rates, noise quickly comes to dominate the signal. To recover accurate expectation values in the presence of noise, a set of error mitigation techniques was applied, including randomized compiling and expectation values rescaling. Pauli Error Reconstruction [34] was applied to learn a stochastic Pauli noise model and a suppression coefficient. A careful study of the error contributions from gates, readout, time-step size and Trotter order was performed, providing a framework to optimize the use of limited quantum resources. By combining these strategies, a 10 times improvement in achievable simulation time was achieved, allowing the study of dynamics up to longer timescales. An example of the simulated dynamics is shown in figure 1.3.

The results of this study were published in the Journal of Plasma Physics [2].

1. Introduction

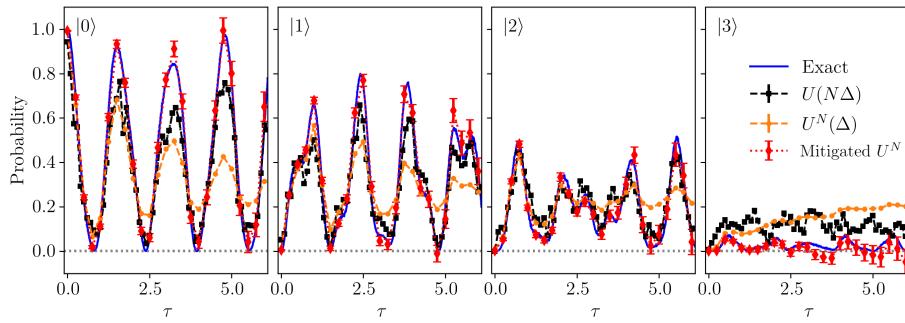


Figure 1.3 – The results of the digital simulation is shown. The blue lines are exact solutions from a classical computer, and the colored symbols with error bars are results from the quantum device with the raw values in orange, and the mitigated values in red. The black points are the result of simulated the exact unitary in a single step on the device.

1.4.3 Simulating electron-plasma waves on a superconducting quantum chip

Despite the impressive increases in circuit depth achieved through error suppression and mitigation, the problem of laser pulse compression proved to be impossible to scale to larger numbers of qubits. This was because the number of gates required to simulate a single timestep scaled unfavorably with the system size, meaning such systems will require fault-tolerant quantum computers to be simulated at scale. Thus, a different class of problems was considered, that of electron-plasma waves and ion-acoustic waves.

A theoretical connection between the well-studied quantum Heisenberg model and plasmas was identified. Specifically, the spectrum of the spin Hamiltonian matches that of electron-plasma waves and ion-acoustic waves, allowing for the study of plasma waves using a larger number of qubits. The spin-chain can be evolved in time using Trotter evolution, requiring only local interactions. Importantly, this results in a circuit which is of a feasible depth for current NISQ devices. Three experiments of interest were identified: First, a many-body Ramsey-type experiment to measure the electron-plasma wave spectrum. Second, the injection of waves into a plasma medium, which could have direct applications in Doppler backscattering, a type of reflectometry technique. Finally, motivated by waves used as actuators in plasmas, the transmission of waves through an inhomogeneous plasma medium was simulated. Simulations of wave propagation in an inhomogeneous plasma are needed to predict where the waves deposit their energy and how the controller will perform in practice. These three cases provide a small-scale demonstration of how quantum computers can be applied to plasma simulation problems of interest, with an efficient mapping to near-term hardware.

Building on the error suppression and optimal control work, a hybrid approach to the circuit was taken which combined digital 1Q rotations with an analog entangling layer. Digital 1Q rotations provided high-fidelity, parameterized gates, while the more challenging entangling operations were implemented by a layer of simultaneous entangling pulses. Rather than attempting to synthesize a particular unitary, the natural FSIM en-

1. Introduction

tangling gates of the device are characterized. To express the target unitary using these native gates, which vary from qubit to qubit, a numerical compilation scheme was leveraged. With the ability to perform high-fidelity timesteps, a high circuit depth became achievable. Error mitigation techniques were adapted for this experiment, leveraging hardware-accelerated randomization and rescaling of expectation values. Combining these techniques, the Rigetti Ankaa-9Q-3 QPU was used to simulate plasma waves up to 15 timesteps, using 60 2Q gate layers. Critically, the techniques and problem are scalable, meaning that further improvement in hardware error rates should allow this experiment to be performed at larger numbers of qubits.

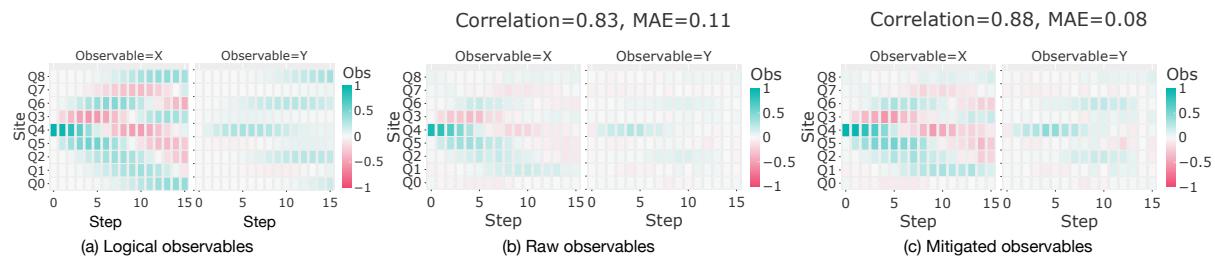


Figure 1.4 – The time evolution of $\langle X \rangle$ and $\langle Y \rangle$ is shown in this many-site Ramsey-type experiment. The time evolution of the complex phase of $\langle X + iY \rangle$ reveals the energies of all the electron-plasma wave eigenmodes. The subplots show the evolution without noise (left), the raw estimated observables (middle), and the error-mitigated results (right).

1.5 Conclusion and outlook

The potential for a quantum advantage in the simulation of plasma physics was investigated both theoretically and experimentally. An optimal control scheme was implemented to enable the synthesis of arbitrary unitaries given a pair of interacting qubits or qutrits. Simultaneously, a set of error suppression and error mitigation techniques were developed on a two-qubit circuit for the simulation of nonlinear optical processes, yielding a dramatic improvement in achievable simulation time. To scale to larger systems sizes, a significant theoretical development mapping electron-plasma waves to spin chains was made, identifying a class of plasma simulation problems that can be efficiently mapped to near-term quantum computers within realistic circuit depths. A scalable experimental protocol was developed which incorporated a digital-analog gate scheme, numerical compilation, and a suite of error mitigation techniques. This supported simulations using 9 qubits and 15 time steps, with experimental results matching well to theory.

While these results represent a great deal of progress in simulation of plasma dynamics using quantum computers, questions remain. Further investigation is required to understand what state-of-the-art classical methods can achieve on this class of problems, in order to assess where a quantum advantage might lie. Methods of targeting non-local interactions would greatly expand the class of problems that could be simulated by quan-

1. Introduction

tum computers. Finally, investigation in richer control schemes, which use multi-qubit interactions could enable denser circuits and simulations.

1.6 Acknowledgments

The following people contributed to the work presented in this report.

Amy F. Brown, Yujin Cho, Jonathan L. DuBois, Bram Evert, Vasily Geyko, Ilon Joseph, Daniel Lidar, Matt Reagor, Eyob A. Sete, Yuan Shi, Bhuvanesh Sundar, Vinay Tripathi.

Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

This chapter summarizes the results of demonstrating optimal control on quantum devices, published in [1].

2.1 Overview

Advanced simulations and calculations on quantum computers require high-fidelity implementations of quantum operations. The universal gateset approach builds complex unitaries from a small set of primitive gates, often resulting in a long gate sequence, which is typically a leading factor in the total accumulated error. Compiling a complex unitary for processors with higher-dimensional logical elements, such as qutrits, exacerbates the accumulated error per unitary, since an even longer gate sequence is required. Optimal control methods promise time- and resource-efficient compact gate sequences and, therefore, higher fidelity. These methods generate pulses that can directly implement any complex unitary on a quantum device. In this work, we demonstrate that any arbitrary qubit and qutrit gate can be realized with high fidelity, which can significantly reduce the length of a gate sequence. Our work promises that the calibration overheads for optimal control gates can be made small enough to enable efficient quantum circuits based on this technique.

2.2 Method

We tested optimal control pulses generated for random unitaries on four superconducting transmon quantum processors, the LLNL Quantum Device and Integration Testbed (QuDIT)'s standard QPU and Rigetti's Ankaa-2, Ankaa-9Q-1, and Aspen-M-3. The QuDIT device has a single transmon made of tantalum on a sapphire substrate that has a long energy decay time [35, 36]. We choose one representative transmon on each of Rigetti's multiqubit chips to focus on in this work. The hardware parameters on the four systems are listed in Table 2.1.

2. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

Table 2.1 – Parameters for the LLNL QuDIT’s standard QPU and the selected qubits on Rigetti’s Ankaa-2, Ankaa-9Q-1, and Aspen-M-3. ω_{ij} indicates the transition frequency from $|i\rangle$ to $|j\rangle$. T_1^{ij} is the energy decay time and T_2^{ij} is the decoherence time in $i - j$ manifold. The T_2^* times were measured with Ramsey oscillation.

Parameters	LLNL QuDIT	Ankaa-2	Ankaa-9Q-1	Aspen-M-3
ω_{01}	3.446 GHz	4.477 GHz	4.526 GHz	3.883 GHz
ω_{12}	3.237 GHz	Not measured	Not measured	Not measured
T_1^{01}	220 μ s	27 μ s	17 μ s	22 μ s
$T_2^{*,01}$	22 μ s	21 μ s	16 μ s	42 μ s
T_1^{12}	145 μ s	Not measured	Not measured	Not measured
$T_2^{*,12}$	25 μ s	Not measured	Not measured	Not measured

In a closed quantum system, the Hamiltonian H of a superconducting transmon in the rotating frame is approximated by

$$H = 0.5\alpha a^\dagger a^\dagger aa + p(t)(a + a^\dagger) + i q(t)(a - a^\dagger), \quad (2.1)$$

up to $\mathcal{O}(a^\dagger a)^2$, where $\alpha = \omega_{12} - \omega_{01}$ is the anharmonicity, a is the lowering operator, and $p(t)$ and $q(t)$ are the control pulses given as time-dependent functions that we optimize. ω_{ij} indicates the transition frequency between $|i\rangle$ and $|j\rangle$. For each target unitary, pulses were obtained either using a Julia open source package, JuQbox (for Rigetti’s devices), or TensorOptimalControl (for LLNL QuDIT’s device).

The pulse lengths were set to 220 ns for the QuDIT qutrit gates, and 40 ns for the Ankaa-2 qubit gates. In all tested cases, we achieved numerical convergence of the pulses with fidelity higher than 99.98 %. Subsequently, we downsampled the pulses to achieve a rate of 1 point per nanosecond for hardware implementation, by selecting 1 point out of every 64 data points. It is important to note that the control pulses do not contain frequency components that exceed 1 GHz, ensuring that no information is lost during the downsampling process.

To achieve the best optimal control gate performance, we calibrate the pulse amplitudes using two scalar factors, γ and σ . The calibrated pulses, \mathcal{C} , are written as:

$$\mathcal{C}(g(f)) = \gamma \left[\mathcal{X}(f < \omega_c) + \sigma \cdot \mathcal{X}(f > \omega_c) \right], \quad (2.2)$$

where $g(f)$ is the frequency domain representation of the control pulses, $p(t)$ and $q(t)$, $\mathcal{X}(\Delta f)$ is the spectral component in the frequency range Δf , and ω_c is the average frequency of ω_{01} and ω_{12} . The amplitude scaling constant, γ , converts pulses calculated in units of the frequency (Hz) to the physically applied voltage by the arbitrary waveform generators. To fine-tune γ , we measure state populations after performing a gate one to ten times, compare them to the predicted trajectories from the Lindblad master



2. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

equation implemented in Python QuTiP package [37, 38], and update γ to minimize the difference between the measured and the predicted evolution. This process allows us to optimize γ for any gate. By adjusting the spectral weight σ between the ω_{01} and ω_{12} components, we compensate for any frequency dependence in the signal chain from the room temperature electronics to the device at 10 mK. The previous method of finding the weight factor constructs a densely-sampled spectral filter around the transition frequencies [39], which could take longer than 10 minutes to measure. In all gates we sampled, the highest spectral components of the pulses were at ω_{01} and ω_{12} , enabling us to instead perform a two-point calibration with a constant weight, σ , applied to the two transition frequencies.

We used a square-root of 0 – 2 swap gate, $sSW02$, as a reference for qutrit gates.

$$sSW02 = \begin{pmatrix} 1/\sqrt{2} & 0 & -i/\sqrt{2} \\ 0 & 1 & 0 \\ -i/\sqrt{2} & 0 & 1/\sqrt{2} \end{pmatrix} \quad (2.3)$$

Using this gate, we adjust γ and σ , until we find the optimal gate performance. In every measurement, we corrected the state-preparation-and-measurement (SPAM) error with an independently-measured confusion matrix. The best calibration parameters for $sSW02$ were then applied to all random gates tested in this work. For qubit gates, we used $RX(\pi/2)$ pulse as a reference and adjusted the amplitude γ to find the optimal calibration parameters.

A simplified calibration procedure can guarantee a high fidelity for any random gate. We fix γ and σ to the calibrated values obtained from $sSW02$, and apply them to the pulses for 300 randomly generated qutrit gates, which were generated from a function, `rand_unitary`, in the QuTiP python package [37, 38]. Figure 2.1 shows measurements of an example randomly generated qutrit gate, G , whose unitary is given by:

$$G = \begin{pmatrix} 0.599 + 0.194i & 0.673 + 0.026i & -0.387 + 0.020i \\ -0.654 + 0.163i & 0.638 - 0.335i & 0.155 - 0.054i \\ 0.387 + 0.027i & 0.164 + 0.011i & 0.887 - 0.191i \end{pmatrix}.$$

Figure 2.1(a-b) shows the reference and the measured process matrices. The magnitude error is less than 0.081 and the average amount of phase error is 0.07π as shown in Fig. 2.1(c-d). By comparing the measured and the reference matrix, we obtained 99.0 % fidelity of the gate G . We repeat the same analysis for 299 different random gates on the QuDIT's QPU. The mean fidelity is 97.9 ± 0.5 % while the highest fidelity is 99.0 %, as presented in Fig. 2.1(e).

This calibration procedure can easily be transferred to a different hardware architecture. To demonstrate this, we follow the same calibration procedure and measure the qubit random gate fidelities on Rigetti's Ankaa-2, Ankaa-9Q-1, and Aspen-M-3. For 100 randomly generated qubit gates, the highest and lowest fidelities are 99.9 % and 96.4 %, respectively. The average fidelities are 99.1 ± 0.4 % for Ankaa-9Q-1 and 99.5 ± 0.3 % for Aspen-M-3 as shown in Fig. 2.2

2. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

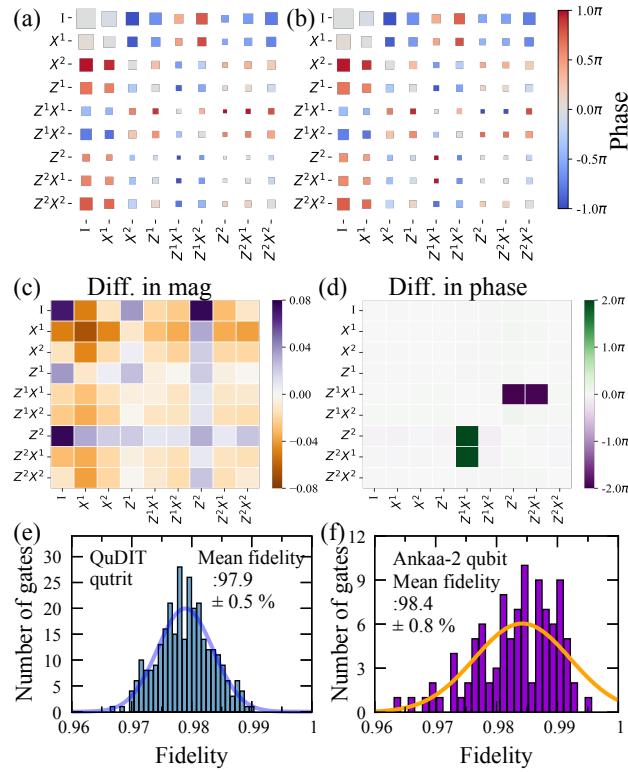


Figure 2.1 – (a) The reference and (b) the measured process matrices of a random gate. The difference between the two matrices is shown in panel (c-d) for magnitude and phase, respectively. (e) The average fidelity of 300 random qutrit gates on the LLNL’s QuDIT platform is $97.9 \pm 0.5\%$. (d) The average fidelity of 100 random qubit gates on Ankaa-2 is $98.4 \pm 0.8\%$.

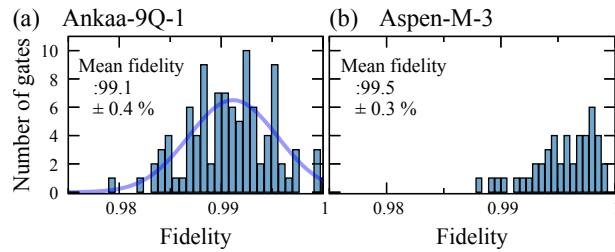


Figure 2.2 – Histogram of fidelities of (a) 100 randomly generated qubit gates on Ankaa-9Q-1 and (b) 46 randomly generated gates on Aspen-M-3.

In the measurements on the tested QPUs, the standard deviations of the fidelities are between 0.3 % and 0.8 %. One possible reason is temporal variations such as fluctuations of the quantum system over time. To measure the temporal fluctuation of fidelity, we monitor the fidelity of the $sSW02$ gate for 3 hours on the QuDIT platform, which is comparable to the duration of the full data set plotted in Fig. 2.1(e). Immediately after the calibration, the measured fidelity was 99.7 %, followed by repeated measurements every 13 – 14 seconds, resulting in 782 measurements. Figure 2.3 shows that the mean

2. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

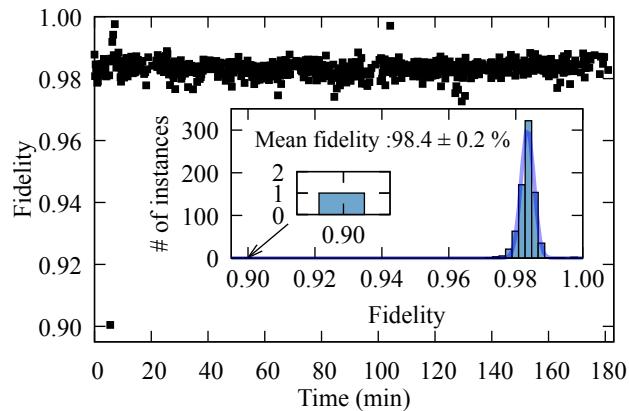


Figure 2.3 – Fidelity of *sSW02* gate changes over time. The mean fidelity is 98.4 % while the highest fidelity is 99.6 %.

fidelity is $98.4 \pm 0.2\%$ with a fidelity drop at one point, as low as 90 %. This low-fidelity is quickly recovered to 98.6 % on the next measurement. The highest fidelity measured is 99.6 %. This 0.2 % fluctuation in the fidelity of a single gate over time partially accounts for the 0.5 % spread in gate fidelities in our random sample (Fig. 2.1(c)). We will discuss the cause of temporal fluctuation and other potential reason for the fidelity span later in Discussion.

QPT provides a thorough information about a gate process. However, QPT is not immune to certain sources of error, such as the SPAM error and the infidelities of operators used for state preparation and projection. Firstly, we correct for the SPAM error by applying a confusion matrix. Secondly, we mitigate the operator infidelities by measuring QPT with gate folding (gQPT). By applying the test gate repeatedly, we selectively amplify errors only in the gate itself, while keeping the errors in the operators unchanged. This allows

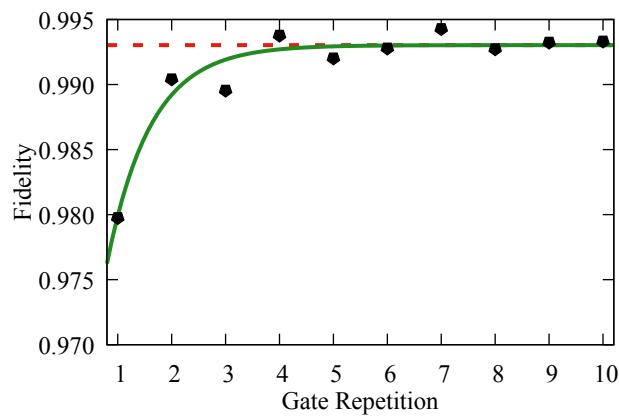


Figure 2.4 – The fidelity of a single random gate was extracted by performing QPT with 1 to 10 applications of the same gate with the spectral weight $\sigma = 1.87$. The fidelity converges to 99.3 %. The green line is a guideline to visualize the convergence of the fidelities.

2. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

us to isolate and quantify the errors specifically associated with the gate under test. By doing so, we can accurately assess the performance of the gate and identify potential areas for improvement. Assuming that errors in each gate operation remains unchanged during the measurement and the operation time is short enough to be insensitive to decoherence, the fidelity can be evaluated by $1/n$ -th root of the measured fidelity at n -th gate folding. In other words, fidelity \mathcal{F} is measured at n -th repetition of gate K , which absorbs the fidelity of the preparation and projection operators, and the adjusted fidelity is $\mathcal{F}^{1/n}$.

Figure 2.4 shows the extracted fidelity of one random gate as a function of the number of gate repetitions. The final gate fidelity is obtained from the exponential fit (green curve), given by $a \cdot \exp(-n/b) + c$ where a , b , and c are the fitting parameters. The parameter c corresponds to the converged fidelity that we extract. We repeat gQPT for 8 different random qutrit gates whose initial fidelity ($n = 1$) is $97.9 \pm 0.3\%$ on average, and overall, the gate fidelity improves and converges to a higher fidelity around 99% . At least 1% of the gate fidelities measured on the QuDIT QPU can be accounted for by the infidelity of the state preparation and projection operators. Similarly, we performed gQPT on 15 randomly selected qubit gates on Rigetti's Ankaa-2 (initial average fidelity: $97.3 \pm 0.6\%$) and achieved average fidelity $99.7 \pm 0.1\%$. To compare gQPT to other existing characterization protocols, we measured the fidelities of these 15 random gates using cross-entropy benchmarking (XEB), which gives us an average fidelity of $98.9 \pm 0.6\%$.

2.3 Discussion

Optimal control can be an efficient tool to implement any high fidelity gates, as we present in this work. Our work shows that we can achieve $98 - 99\%$ average gate fidelity for randomly-generated gates with fast one-time calibration that can be applied to any random gate we want to produce.

To measure the true fidelities of test gates apart from SPAM error and operator infidelities, we selectively amplified errors in the test gate by gate folding and extracted the fidelity from an exponential fit (Fig. 2.4). QPT and gQPT are valid in the presence of small coherent errors, which are explicitly identified in the measured process matrix and contribute to the reported infidelities. When the coherent errors in the operators dominate the error in the tested gate, gQPT minimizes the impact of operator infidelities and the estimated fidelity rises with the gate application and asymptotes towards a fixed value, as we saw in our reported experiments. In our measurements, coherent errors in all tested gates are small enough for gQPT to be valid for up to 10 gate repetitions; the lowest gate fidelity measured by conventional QPT is 96.7% (qutrit on QuDIT) and 96.3% (qubit on Ankaa-2). The extracted gate fidelity from gQPT increases by about 1% over conventional QPT and is about 0.8% higher than XEB measurements. We attribute the difference to the operator infidelities in conventional QPT and XEB measurements. XEB is inherently free from measurement errors, but not from operator infidelities. Single

2. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

qubit primitive gates can be calibrated to very high fidelities, but small infidelities can accumulate at each application, which can contribute to slightly lower fidelities of XEB measurements. The measured fidelity with gQPT is an approximation to the actual gate fidelity as it does not capture other types of errors, such as statistical errors (shot noise) which typically follows a Gaussian distribution.

During the experiments, we have observed fluctuation of a gate fidelity over time as shown in Fig. 2.3. Quantum systems, in general, exhibit fluctuations and drift on the time scale of minutes, which can be attributed to various noise sources. For instance, inherent defects in quantum devices, such as residues from the fabrication process or adsorbents from the air, can cause resonant coupling with the qubits, leading to shifts in qubit frequencies and a sharp decrease in T_1 [40]. Additionally, cosmic rays can generate quasiparticles that temporarily reduce T_1 , which affects the quantum system on the order of milliseconds. Interaction with the environment can cause a long-term drift over a few hours. Due to the need for control, quantum devices cannot be perfectly isolated from the environment. Even though the coupling is weak, small environmental changes, like fluctuations in room temperature can introduce noise into the device. Similarly, instability in room temperature electronics can introduce control errors. For instance, the amplitude of a π pulse needs occasional recalibration to ensure optimal performance. Reducing noise sources and mitigating their impact on qubit performance is an active area of research. The calibration parameters are stable for several hours and require a fine-tuning of the amplitude at the 1 %-level after two weeks due to drift in the lab environment and the quantum system.

In addition, the standard deviations of the measured fidelities are 0.3 – 0.8 %, which is larger than expected from temporal fluctuation. This fidelity range could be related to the choice of universal parameters. Different qutrit gates have different ratio of the spectral components between the 0 – 1 and 1 – 2 transitions, which could lead to slightly adjusted spectral weight σ . For example, we optimized the spectral weight for the specific gate we present in Fig. 2.4. For this gate, the highest fidelity was at $\sigma = 1.87$, instead of the $\sigma = 1.8$ for the reference gate. This suggests that adding a parameter that depends on the weight of spectral component may help implementing higher gate fidelities with narrower standard deviation. On the other hand, arbitrary qubit gates are independent of σ , because the pulse has only one frequency component. In this case, the fidelities of arbitrary qubit gates could be improved by frequent tuning of the amplitude scaling constant γ . Quantum hardware requires regular tuning of quantum gates to ensure its best performance by optimizing pulse amplitude. If the amplitude from control electronics is unstable, the system would require more frequent tuning of the pulse amplitude.

We have tested arbitrary qutrit gates using optimal control on Rigetti's Aspen-M-3 device. However, we achieved only $85.6 \pm 3.6 \%$, which is more than 10 % lower than what we measured on LLNL's QPU. This result indicates that the performance of arbitrary qutrit gates implemented with optimal control is system-dependent. Multi-qubit chips typically exhibit higher levels of noise compared to single qubit devices, primarily due to additional noise sources like flux or interactions between qubits. These noise sources

2. Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits

are not static and can change over time. In order to determine the frequencies required for the 0-1 and 1-2 transitions in the Hamiltonian model used for optimal control, we relied on Rigetti's calibration data, which is updated every 6 hours. However, if the qubit frequencies experience significant drift within this 6-hour time frame, the accuracy of the Hamiltonian model decreases, leading to phase errors accumulating during pulse execution. Towards the end of the 6-hour calibration block, we noticed increasing coherent errors when using the Rigetti's pre-defined native gates. This suggests that the amplitude factor for the optimal control pulses also needs to be recalibrated. Unfortunately, due to limited resources and accessibility to the Rigetti device, we were unable to thoroughly investigate all potential sources of error.

One future direction is to apply this method to more robust optimal control gates. When we monitor the fidelity of one gate over a few hours, the fidelity drops by 8% at one time and the standard deviation is 0.2% as presented in Fig. 2.3. To generate robust pulses for a target unitary over time that is more stable, it would be useful to build a model that captures the time-dependence of the system, such as T_1 , T_2 , and qubit frequency fluctuation that is typically in the order of a few kHz [41]. Another direction would be to explore systematic way to achieve fast control of an arbitrary gate. To achieve the shortest gate time, there are a few challenges to overcome. When a gate becomes shorter, it tends to have higher amplitudes, which can unintentionally drive higher energy excitation. In addition, the pulse length is often limited by the clock cycle of the arbitrary waveform generator, typically 4 ns, limiting our ability to fully explore the dynamic range of the pulse length. Lastly, we plan to expand this calibration method to multi qubit and qutrit entangling gates to achieve high performance entangling operations.

2.4 Conclusion

In this work, we experimentally demonstrate that optimal control technique can prepare any random quantum logic gate with minimal calibration at high fidelities, opening the door to greater adaptation of this technique. Our calibration procedure is applicable to different hardware architectures, showing that the optimal control is a practical and promising direction for optimized quantum circuits.

Implementing custom gates at the pulse-level enables us to operate quantum simulations and algorithms faster with higher fidelities. For example, in Quantum Fourier Transform [42, 43] or variational quantum eigensolver (VQE) [44, 45, 46, 47], a sequence of fixed gates can be replaced with an optimal control pulse to reduce the operation time and the overall gate count in the circuit. Similar ideas have been suggested to use parametrized pulses as an ansatz for higher fidelity VQE calculation [48].

3

Simulating nonlinear optical processes on a superconducting quantum device

Simulating plasma physics on quantum computers is difficult because most problems of interest are nonlinear, but quantum computers are not naturally suitable for nonlinear operations. In weakly nonlinear regimes, plasma problems can be modeled as wave-wave interactions. In this paper, we develop a quantization approach to convert nonlinear wave-wave interaction problems to Hamiltonian simulation problems. We demonstrate our approach using two qubits on a superconducting device. Unlike a photonic device, a superconducting device does not naturally have the desired interactions in its native Hamiltonian. Nevertheless, Hamiltonian simulations can still be performed by decomposing required unitary operations into native gates. To improve experimental results, we employ a range of error mitigation techniques. Apart from readout error mitigation, we use randomized compilation to transform undiagnosed coherent errors into well-behaved stochastic Pauli channels. Moreover, to compensate for stochastic noise, we rescale exponentially decaying probability amplitudes using rates measured from cycle benchmarking. We carefully consider how different choices of product-formula algorithms affect the overall error and show how a trade-off can be made to best utilize limited quantum resources. This study provides an example of how plasma problems may be solved on near-term quantum computing platforms.

3.1 Classical model of laser pulse compression

An important class of wave-wave interactions in plasma physics are laser-plasma interactions. As an example, we consider a scenario where a plasma is used for laser pulse compression [49], during which the intensity of a seed laser pulse is amplified while its duration is shortened. Classically, the laser amplification is often treated as a parametric process, where the signal and idler waves grow by consuming a pump wave. When the pump energy is being replenished, or when the pump energy dominates, one may approximate the pump amplitude a_1 as a constant, in which case the seed amplitude a_2 and the idler amplitude a_3 grow exponentially. However, when the pump amplitude is not held constant, the three-wave nature of the underlying interaction becomes apparent.

3. Simulating nonlinear optical processes on a superconducting quantum device

The interaction is often described by the three-wave equations [50]

$$d_t a_1 = g a_2 a_3, \quad (3.1)$$

$$d_t a_2 = -g^* a_1 a_3^\dagger, \quad (3.2)$$

$$d_t a_3 = -g^* a_1 a_2^\dagger, \quad (3.3)$$

where d_t is the advective derivative, g is the coupling coefficient, g^* is its complex conjugate, and a^\dagger denotes the complex conjugate of a in the classical model. The advective derivative is specific for each wave and is defined as $d_t = \partial_t + \mathbf{v}_g \cdot \nabla + \mu$, where $\mathbf{v}_g = \partial \omega / \partial \mathbf{k}$ is the group velocity and μ is the damping rate of the wave. The complex-valued amplitude a is the slowly varying envelope of the classical wave. The amplitude is normalized such that $n = |a|^2$ is proportional to the wave action density, which is proportional to the number of photons in the wave.

At later stages of pulse compression, the intensity of the seed far exceeds the pump. The large a_2 induces an additional relativistic nonlinearity. The nonlinearity originates from the fact that, in plasmas, photons are massive particles due to their interactions with free charges. In unmagnetized plasmas, photons satisfy the dispersion relation $\omega^2 = \omega_p^2 + c^2 k^2$, where the photon mass can be identified with the plasma frequency $\omega_p = (e^2 n_e / \epsilon_0 m_e)^{1/2}$. Here, e is the electron charge and n_e is the electron density. Because electrons oscillate in the laser's electric field, the effective electron mass m_e is replaced by γm_e when the electron quiver speed v_q becomes comparable to the speed of light c , where $\gamma = 1 / \sqrt{1 - v_q^2 / c^2}$. As the seed pulse propagates, at places where the pulse is more intense, the photon mass $\omega_p \propto \gamma^{-1/2}$ becomes smaller. A smaller ω_p leads to a larger k at a fixed ω , which means a larger group velocity $\mathbf{v}_g = c^2 \mathbf{k} / \omega$. Consequently, the more intense part of a_2 moves at a higher group velocity. If the envelope of a_2 has initial modulations, then they will pile up and grow. This process is known as relativistic modulational instability. The equation that describes the modulational instability of the seed pulse is

$$d_t a_2 = i R a_2^\dagger a_2 a_2, \quad (3.4)$$

where d_t is again the advective derivative and $R = \omega_p^2 / (8\omega)$ is the coupling coefficient. For relativistic modulational instability, $R > 0$ is a real number, which means $iR|a_2|^2$ is purely imaginary. The above equation thus modulates the phase of the complex a_2 in such a way that a larger $|a_2|$ leads to a faster phase evolution. Since the three-wave interaction is a phase sensitive process, the modulational instability spoils the amplification process by introducing a phase mismatch.

3.2 Quantum model of laser pulse compression

In the classical model, the amplitude a is a complex-valued function, and $n = |a|^2$ is proportional to the number of photons. This setup naturally admits canonical quantization for bosonic quantum fields $[a_i(\mathbf{x}), a_j^\dagger(\mathbf{y})] = \delta_{ij} (2\pi)^3 \delta^{(3)}(\mathbf{x} - \mathbf{y})$, where the indices

3. Simulating nonlinear optical processes on a superconducting quantum device

$i, j = 1, 2, 3$, and the operators have spatial dependencies. Since we will later implement the model on quantum hardware, which has a limited number of qubits, in this paper we will focus on the temporal problem with no spatial dependence. In this case, when damping is negligible, the advective derivative $d_t \rightarrow \partial_t$ is reduced to a partial derivative in time, and the operators satisfy the canonical quantization conditions

$$[a_i, a_j^\dagger] = \delta_{ij}. \quad (3.5)$$

The quantization promotes normalized amplitudes to creation and annihilation operators, and the Kronecker delta distinguishes the three types of waves in the system. For each wave type, the number operator is $n_i = a_i^\dagger a_i$. The eigenstates of n_i are the Fock states $|m_i\rangle$. Since we have three types of waves, it is convenient to abbreviate the tensor-product state $|m_1\rangle \otimes |m_2\rangle \otimes |m_3\rangle$ as $|m_1, m_2, m_3\rangle$. This number basis is natural for the quantized problem.

While the Schrödinger picture $i\partial_t|\psi\rangle = H|\psi\rangle$ is more convenient for quantum simulations, the connection between the quantum and classical models is more transparent in the Heisenberg picture $d_t a = i[H, a]$. For three-wave interactions, Eqs. (3.1) are the Heisenberg equations from the cubic Hamiltonian

$$H_T = ig a_1^\dagger a_2 a_3 - ig^* a_1 a_2^\dagger a_3^\dagger. \quad (3.6)$$

The first term of H_T annihilates a seed and an idler photon to create a pump photon, while the second term of H_T is the reverse process where a pump photon decays into a seed and an idler photon. Although the Heisenberg equations for a_j are formally identical to the classical three-wave equations, the difference between the quantum and classical systems become apparent when one calculates higher order cumulants. For example, because a_j and a_j^\dagger do not commute, the Heisenberg equation for n_i is different from its classical counterpart [51]. Similarly, for the four-wave interaction, Eq. (3.4) is the Heisenberg equation from the quartic Hamiltonian

$$H_F = -\frac{R}{2} a_2^\dagger a_2^\dagger a_2 a_2, \quad (3.7)$$

which is also known as the self-Kerr nonlinearity in the quantum literature. Since $R > 0$ for the modulational instability, the negative sign in H_F means that photons tend to condense together, which leads to a lower energy of the system.

To use quantum Hamiltonian simulations to solve the quantized wave-wave interaction problems, we focus on the Schrödinger picture and use a basis that respects action conservation. In classical wave-wave interactions, $S_2 = n_1 + n_2$ and $S_3 = n_1 + n_3$ are known as the conserved wave actions. In the quantized model, $[H_T, S_2] = [H_T, S_3] = 0$, and H_F also commutes with S_2 and S_3 . Therefore, it is convenient to use eigenstates of S_2 and S_3 as the computational basis, which we shall call the action basis. For the laser pulse compression problem, since we are primarily interested in the seed wave a_2 , we label the action basis by

$$|\phi_j^{s_2, s_3}\rangle = |s_2 - j, j, s_3 - s_2 + j\rangle, \quad (3.8)$$

3. Simulating nonlinear optical processes on a superconducting quantum device

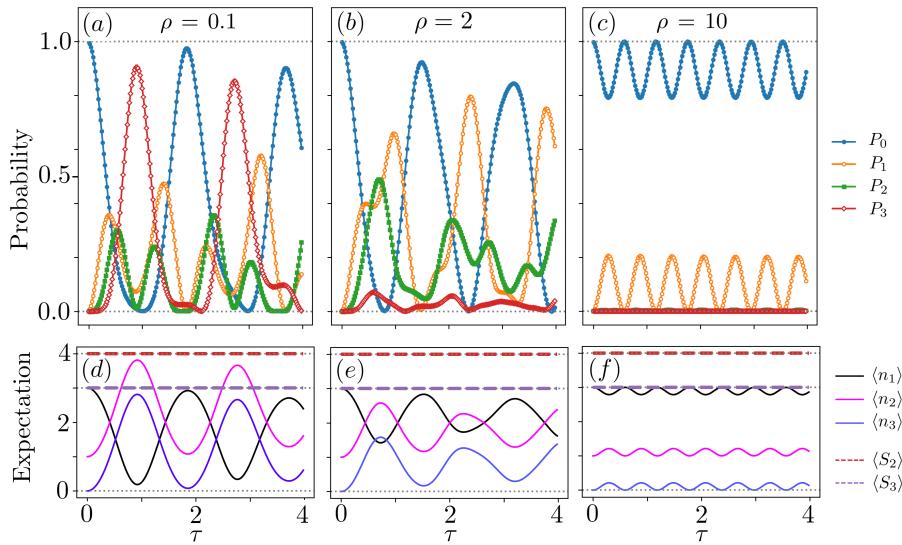


Figure 3.1 – Exact dynamics of mixed three- and four-wave interaction problems in a $D = 4$ dimensional Hilbert space with constants of motion $s_2 = 4$ and $s_3 = 3$. Starting from the ground state, the probability amplitudes c are evolved in time, and the occupation probabilities $P_l = |c_l|^2$ (a)-(c), as well as the expected quanta in the three waves $\langle n_i \rangle$ (d)-(f) are computed on a classical computer. When $\rho = R/|g| = 0.1$, three-wave interaction dominates; when $\rho = 2$, three- and four-wave interactions compete; when $\rho = 10$, four-wave interaction dominates.

where j is the number of photons in the seed wave. Because the Hamiltonian is block diagonal in the action basis, we can perform Hamiltonian simulations separately in each (s_2, s_3) subspace. In other words, the infinite dimensional Hilbert space can be decomposed as a direct sum of finite dimensional subspaces, where each subspace is labeled by a pair of quantum numbers (s_2, s_3) .

3.3 Exact quantum dynamics

The exact dynamics involves two fundamental frequency scales g and R , from H_T and H_F , respectively [Eqs. (3.6) and (3.7)]. When $g = 0$, the dynamics is trivial because H_F is diagonal: Under the influence of H_F alone, the occupation of $|\phi_j\rangle$ remains unchanged, and the dynamics is a pure phase precession. To change occupation numbers, a nonzero g is needed. Hence, for nontrivial dynamics, we can always normalize time by $\tau = |g|t$ and normalize the four-wave coupling by $\rho = R/|g|$. Since the Hamiltonian is time independent, the exact dynamics is described by the unitary evolution operator $U(\tau) = \exp(-iH\tau)$, which can be obtained by direct diagonalization and exponentiation, at least for small problem sizes. For size $D = 4$, the exact behaviors of three examples are shown in Fig. 3.1. In all three examples, the coupling phase $\theta = 0$ and the constants of motion are $s_2 = 4$ and $s_3 = 3$, which means $j_{\min} = 1$ and $k = l + 1$, so $|\phi_k\rangle = |3 - l, 1 + l, l\rangle$. All examples start from the ground state of the computational basis.

3. Simulating nonlinear optical processes on a superconducting quantum device

The dynamics is controlled by the dimensionless parameter ρ . When $\rho \ll 1$, as shown in Fig. 3.1(a) and 3.1(d), three-wave interaction dominates, which causes population transfer between the three waves. In this case, the much weaker four-wave interaction slowly accumulate phase mismatches that reduces the efficiency of population transfer, which is manifested by the decreasing oscillation amplitudes of $\langle n \rangle$ in Fig. 3.1(d). Here, the expected number of quanta in the three waves are calculated from the probability amplitudes c by $\langle n_1 \rangle = \sum_l (s_2 - j) |c_l|^2$, $\langle n_2 \rangle = \sum_l j |c_l|^2$, and $\langle n_3 \rangle = \sum_l (s_3 - s_2 + j) |c_l|^2$, where $j = j_{\min} + l$, and the summation is over $l = 0 \dots, D - 1$. From the above equations, it is clear that $\langle S_2 \rangle = s_2$ and $\langle S_3 \rangle = s_3$ are exact constants of motion, as marked by horizontal dashed lines in the lower panels of Fig. 3.1. In the opposite limit $\rho \gg 1$, as shown in Fig. 3.1(c) and 3.1(f), four-wave interaction dominates. In this case, the phases of different $|\phi_j\rangle$ precess at drastically different rates, which inhibits population transfer because three-wave interaction requires phase matching. In this example ($\rho = 10$), the $0 \leftrightarrow 1$ transfer is strongly suppressed, while the $1 \leftrightarrow 2$ transfer, which accumulate phase mismatch at a greater rate, becomes nearly impossible. Finally, in the intermediate case $\rho \sim 1$, as shown in Fig. 3.1(b) and 3.1(e), three- and four-wave interactions compete, and the dynamics is more complicated. While four-wave interactions generate phase mismatches that suppress population transfer, three-wave interactions change the populations and affect how the phases are weighted. The intermediate cases are where simulations are most needed for predicting the behavior of the system.

3.4 Implementing exact dynamics with error mitigation

We perform two-qubit experiments on Rigetti's Aspen-M-3 processor [52], which is a superconducting device with multiple transmon qubits at a fixed topology with hardwired qubit-qubit couplings. The device is routinely calibrated to support single-qubit gates, as well as two-qubit gates like CZ and parametric XY(θ) gates. Each experiment is specified as a sequence of unitary operations, and each 4-by-4 unitary matrix is decomposed using Cartan decomposition into at most three two-qubit SQISW gates, sandwiched between single-qubit gates [53]. The total gate sequence is executed on the hardware with the device initialized in the ground state. At the end of the gate sequence, the states of the two qubits are measured. The whole process of an experiment takes a few microseconds to run on hardware, with the overall time being dominated by a passive reset delay. We repeat each experiment for $M = 50,000$ times to accumulate statistics for the final states, so that the shot noise, which scales as $O(1/\sqrt{M}) \sim 0.4\%$, is small compared to other sources of errors. At the end of M repeated experiments, we obtain a single data point along the time history of the evolution. Because projective measurement destroys quantum states, to obtain the next point along the time history, the simulation has to restart from the beginning in the form of a different experiment, which has its own sequence of unitary operations and is repeated another M times.

As the first test of the quantum device, we use it to enact the exact unitary operator. In Fig. 3.2, the solid blue lines are the exact occupation probabilities of the four states in

3. Simulating nonlinear optical processes on a superconducting quantum device

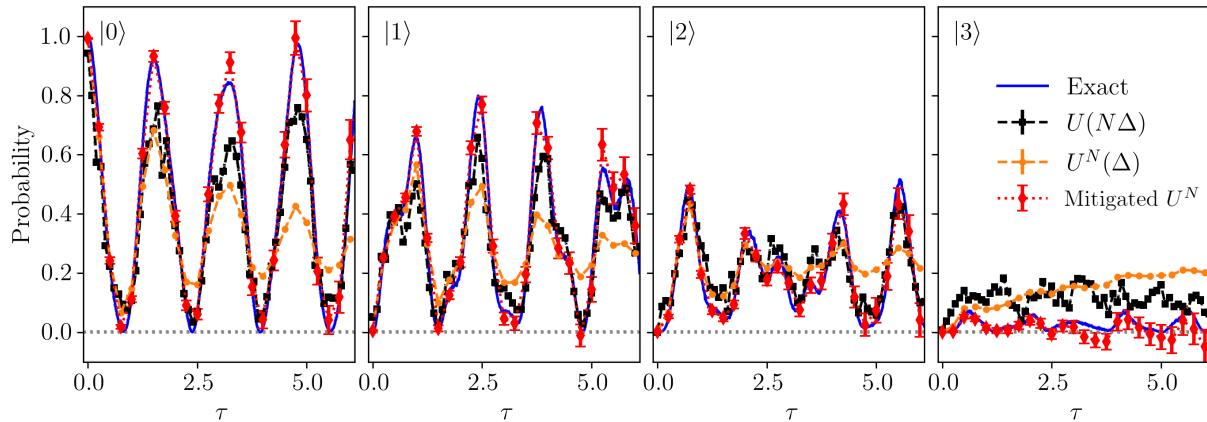


Figure 3.2 – Occupation probabilities in a test problem with parameters $\rho = 2$, $\theta = 0$, $s_2 = 4$, and $s_3 = 3$. The blue lines are exact solutions from a classical computer, and the colored symbols with error bars are results from the quantum device. When asked to enact the final unitary (black), the device performance is acceptable but not ideal. However, when asked to perform time evolution (orange), results on the device degrade to noise level after a few oscillations. The results are significantly improved using error mitigation techniques (red), but the error bars grow exponentially.

our computational basis, which are computed using exact exponentiation on classical computers. The test problem uses parameters $\rho = 2$, $\theta = 0$, $s_2 = 4$, and $s_3 = 3$, which are identical to the middle panel of Fig. 3.1. The exact solutions serve as references for results on the quantum device. This first test is the simplest task that a quantum hardware can perform: For each time $\tau = N\Delta$, we compute the unitary exactly on a classical computer. The sequence of unitary operations for this experiment is thus constituted of just a single unitary, $U(N\Delta)$, and the results are shown in Fig. 3.2 as the black dashed lines. As can be seen from the figure, even when enacting a single dense unitary on the device, the fidelity is far from perfect. In this test, because the gate sequence is so short, decoherence is not a leading cause of infidelity. Instead, most infidelity comes from coherent gate errors, in the sense that each gate realizes a slightly different unitary than what is intended. In this simplest test, another source of error is readout, for which we have already corrected using an iterative Bayesian unfolding technique [54].

As the second test, we perform time evolution using the exact unitary $U(\Delta)$, and the results are shown by the orange dashed lines in Fig. 3.2. In this set of experiments, $U(\Delta)$ is compiled to native gates, and the gate sequence is repeated N times to enact $U^N(\Delta)$. Because of the repetition, as $\tau = N\Delta$ increases, the gate depth increases linearly. The accumulation of errors leads to a degradation of fidelity, as can be seen from Fig. 3.2. The oscillation amplitudes decrease and $p(\tau)$ deviates further from the true solution as τ increases. At even larger τ values, the quantum states become fully scrambled, so $p \rightarrow 1/4$ approaches the fully mixed value for the four quantum states. Because $U^N(\Delta)$ has a larger depth than $U(N\Delta)$, the device performs worse in this test (orange lines) than in the previous test (black lines) as expected. The $U^N(\Delta)$ results improve noticeably from

3. Simulating nonlinear optical processes on a superconducting quantum device

[51] primarily because of the SQISW gate, which has significantly shorter duration and higher fidelity than two-qubit gate used in our previous work.

We need to improve the results for the exact unitary before moving on to the next test. The dominant source of error on the quantum processor are two-qubit (2Q) gate errors, which are typically an order of magnitude larger than single-qubit (1Q) gate errors. On the Aspen architecture, this is not only due to a significantly longer 2Q gate time, but also because activating the 2Q gate requires tuning one of the qubits away from its optimal operating point so the qubit becomes more sensitive to flux noise [55], which leads to a higher dephasing rate for the qubit. We take a multi-pronged approach to minimizing 2Q gate errors. First, the Aspen-M-3 chip used in our experiment has ~ 200 calibrated 2Q gates available. We require only one of these for this circuit and are thus able to select high-performing candidates based on the reported fidelities. Secondly, the Aspen chip offers both native CZ and $XY(\theta)$ gates, thus providing a choice of how to express our problem unitary. We observe that the XY family is particularly expressive [56], allowing the expression of our target unitary using two $XY(\pi/2)$ gates and single-qubit gates [53]. We find that by using SQISW as our native gate, we can cut the gate duration in half and reducing the 2Q gate error by around 40 percent. Our native 2Q gate is thus a 64-ns SQISW which is combined with single-qubit rotations to produce highly expressive native cycles.

To mitigate coherent errors, we convert them into stochastic Pauli errors using a random compilation technique [57, 58]. The technique exploits the fact that the decomposition of a target unitary into elementary gates is not unique. By inserting random single-qubit gates that can be commuted across the two-qubit gate, we generate equivalent gate sequences that are Pauli twirled. If the Pauli twirling gates are chosen independently and if hardware Pauli errors are also independent, then random compilation transforms any gate errors into stochastic Pauli errors. In other words, suppose the errors of a quantum channel, when represented by the Pauli-transfer matrix, have off-diagonal components before twirling. Then, after twirling, the errors become purely diagonal, which means coherent interference of errors is removed. Because our native two-qubit gate is a non-Clifford gate, we cannot apply full Pauli twirling. Rather, we use a pseudo-twirling technique which tailors a smaller subset of coherent errors using the group of single-qubit rotations which can be successfully inverted by 1Q gates. This twirling group is less powerful than Pauli twirling, but still tailors the noise effectively in most situations. The twirling is performed using the TrueQ software library [59]. In our experiments, we construct the logical circuit and compute 50 random compilations. Each compilation has an identical pulse schedule, and thus an identical noise model. The randomization of 1Q gates is performed by updating angles of our virtual Z gates. Such updates can be made with high efficiency, allowing a large number of randomizations of the circuit to be executed in quick succession.

The final step of error mitigation is to compensate for the suppression of observables using a rescaling technique [60]. After twirling of a unitary operation U , the noise channel becomes approximately $\mathcal{E}(\rho) = \sum \lambda_P P \rho P^\dagger$, where the summation is over all tensor products of 1Q Pauli operators P . The coefficient P , called the Pauli decay constant,



3. Simulating nonlinear optical processes on a superconducting quantum device

is specific to the Pauli operator P but is independent of the unitary U that is being performed. Because the λ_P 's are bounded by their mean $\bar{\lambda}$ as $2\bar{\lambda} - 1 \leq \lambda_P \leq 1$, we use a single $\bar{\lambda}$ value to correct for all Pauli errors. This approximation becomes exact when the errors are fully depolarizing, which means that all Pauli channels decay in the same way. In this case, the measured expectation value \tilde{E} for any E of interest is given by a simple rescaling $\tilde{E} = \bar{\lambda}E$, because the length of the Bloch vector, which measures the purity of the state, shrinks by $\bar{\lambda}$. For example, in our two-qubit problem, the expected occupation of a state beyond its fully mixed value is $(\tilde{p} - \frac{1}{4}) = \bar{\lambda}(p - \frac{1}{4})$, where p is the occupation probability for a pure state after one unitary operation. Then, after N unitary operations, we can purify the probability by a rescaling $p = \frac{1}{4} + \bar{\lambda}^{-N}(\tilde{p} - \frac{1}{4})$. In other words, after measuring the probability of a state, we subtract the noise $\frac{1}{4}$ and amplify the remaining signal exponentially by a rescaling factor $(1/\bar{\lambda})^N$. Notice that while amplifying the signal, this purification procedure also amplifies statistical error bars exponentially.

After performing the above error mitigation steps, the hardware results for our test problem are shown in Fig. 3.2 by the dotted red lines. The mitigated results of the second test now closely tracks the exact solutions, and performs even better than the first test (black lines), which does not use any mitigation. While the mitigation significantly improves the signals, without noticeable increasing the hardware overhead, the price we pay is exponentially growing error bars. At even larger simulation depth, the error bars will become comparable to the signals, beyond which the simulations need to stop.

3.5 Testing product formulas and optimal use of quantum resources

With sufficient simulation depth, we can now test the next level of quantum simulations, without assuming that the exact unitary is known. For given problem parameters, we compile unitary matrices $U_T(\tau) = \exp(-iH_T\tau)$ and $U_F(\tau) = \exp(-iH_F\tau)$ to native gates using a Cartan decomposition. Then, we use U_T and U_F to approximate the exact U . With error mitigation techniques, we are able to run experiments on the quantum device for up to about two hundred 2Q gates. The gate depth is deep enough that we can begin to compare results of different product formula algorithms.

With a fixed gate budget, because lower order algorithms require fewer unitary operations per step, they can afford to use smaller time step sizes. In contrast, higher order algorithms require significantly more unitary operations per step, and thus can only afford to use a much larger Δ for a fixed total gate depth. It is worth emphasizing that product formulas indeed become more accurate at higher orders, provided that the time step size Δ is fixed. In our tests, higher order algorithms perform worse because Δ is changed, such that the total gate depth does not exceed what is viable on the quantum device. An analogy here is the run time on classical computers. While higher order algorithms are more accurate at a fixed resolution, they require more operations and therefore longer run time. When given a fixed run time, one is forced to use a coarser

3. Simulating nonlinear optical processes on a superconducting quantum device

resolution, in which case higher order algorithms may perform worse than lower order algorithms.

On current quantum devices, which do not yet have operational error correction, the maximum gate depth is limited. To make the best use of the limited quantum resources, we can adjust the choice of algorithms and resolutions for a given problem. In our case, the pulse compression problem seeks to determine the final seed laser intensity at the end of the interactions. The final time is set, for example, by the duration of laser pulses or the time to traverse the size of the mediating plasma. As a test problem, we fix $\tau_f = 1$ with parameter values $\rho = 4$, $\theta = 0$, $s_2 = 3$, and $s_3 = 3$. We perform Hamiltonian simulation on the quantum device using product formulas to evolve quantum states, and the measured occupation probabilities are post-processed to compute the expectation values of the three waves. We measure the error of the simulation by $\epsilon = \{\frac{1}{N} \sum_{k=1}^N [n(k\Delta) - \langle n(k\Delta) \rangle]^2\}^{1/2}$, where n is the exact result on a classical computer and $\langle n \rangle$ is the expectation value obtained from the quantum hardware. Notice that the error of $\langle n \rangle$ is different from, albeit correlated with, the errors in the unitary U . The unitary error, which is also known as the process infidelity, gives a more complete characterization of the hardware performance. But the expectation-value error is of more interest to the pulse compression problem: it is the same type of error that would typically be determined using a classical algorithm, and is much easier to measure than full process tomography in experiments.

The overall error receives contributions from two fundamental sources. First, algorithmic errors often arise when simulating a dynamical system with knowledge of only the exact solutions of its noncommuting subsystems. Algorithmic errors are unavoidable even on classical computers. In our test problem, we assume that the separate three- and four-wave unitary can be implemented exactly, and then use product formulas to approximate the total unitary. In this case, any finite time step size introduces a discretization error ϵ_Δ , which can be reduced either by using higher order formulas at fixed Δ , or by using the formula at a fixed order but with decreasing Δ . When using the Suzuki formula, demanding errors to scale as Δ^{q+1} per step requires M_q operations, which grows exponentially with q . On the other hand, to reach a target final time τ_f , the number of steps $N = \tau_f/\Delta$ increases only linearly when decreasing Δ . The total algorithmic error $\epsilon_1 = O(\tau_f \Delta^q)$ can in principle be made arbitrarily small by increasing q and decreasing Δ . However, in practice, given a limited run time, finite algorithm precision must be chosen. The trade-off between using a larger q versus a smaller Δ is strongly influenced by the second fundamental source of error: the hardware error. We measure hardware error by ϵ_Q , the error per unitary operation, which is analogous to round-off errors on classical computers. On future error corrected quantum computers, it will be possible to suppress ϵ_Q to arbitrarily small values. However, on current noisy devices, with only error mitigation rather than error correction, ϵ_Q is substantial. Using randomized compilation, we transform coherent errors into random Pauli errors, which contribute to depolarizing noise together with intrinsic quantum decoherence. After error mitigation, the error for one operation becomes independent from the previous operation, so the total hardware error $\epsilon_2 = O(N M_q \epsilon_Q)$ accumulates linearly with the number of operations in the worst-

3. Simulating nonlinear optical processes on a superconducting quantum device

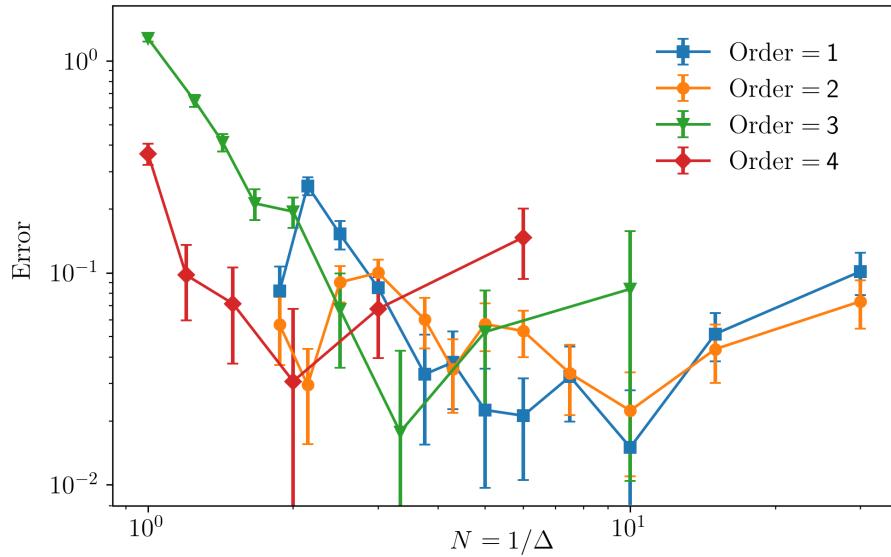


Figure 3.3 – For a given problem, minimum error is obtained at the trade-off between algorithmic errors and hardware errors. All test problems use common parameters $\rho = 4$, $\theta = 0$, $s_2 = 3$, and $s_3 = 3$. The targeted final time $\tau_f = N\Delta = 1$ is fixed, so a finer resolution Δ requires more steps N , which means algorithmic errors decrease with N at the expense of accumulating more hardware errors. An optimal resolution exist, where the overall error is minimized. At higher order, the optimal N shifts towards lower resolution, and the minimum error does not improve with the algorithm order.

case scenario. Because the algorithmic error ϵ_1 is independent of the hardware error ϵ_2 , the overall error is $\epsilon = \epsilon_1 + \epsilon_2$. Notice that ϵ_1 decreases with N , whereas ϵ_2 increases with N , so there is an optimal resolution Δ at which ϵ is minimized.

The trade-off between hardware and algorithmic errors is demonstrated by a suite of experiments, whose results are shown in Fig. 3.3. We test the four product formulas using a common test problem, whose parameters are $\rho = 4$, $\theta = 0$, $s_2 = 3$, $s_3 = 3$, and $\tau_f = 1$. For each order of the product formula, the overall error first decreases with N due to the reduction of algorithmic errors at finer resolution. However, when N exceeds an optimal value N_* , the error starts to increase due to the accumulation of hardware errors. At small N , the resolution is too coarse to resolve the dynamics, so the errors are $O(1)$, which are comparable to the signals. If N_* had been larger, one would expect to see that ϵ for higher order algorithms is smaller and decreases at a steeper slope. In our tests, because N_* is not large enough, such a behavior is not clearly observed. At large N , where the accumulation of hardware errors dominates, ϵ increases roughly linearly with N for all orders. Because higher order algorithms use more operations per step M_q , higher order curves reside above lower order ones in the log-log plots of ϵ , except between orders 1 and 2. Notice that although $M_1 = 2$ and $M_2 = 3$, after merging adjacent unitary operations of the same type, the first-order sequence has $2N$ operations, while the second-order sequence has $2N + 1$ operations, which is only slightly larger. In our test, higher order algorithms do not perform better than a first-order algorithm at their respective optimal N_* . However, this may change on future quantum devices where ϵ_Q

3. Simulating nonlinear optical processes on a superconducting quantum device

becomes even smaller. Similar conclusions have recently been reached for simulating the transverse-field Ising model and the XY model on noisy quantum computers [61].

4

Simulating electron-plasma waves on a superconducting quantum chip

4.1 Introduction

Simulations of three-wave and four-wave interactions [62, 2] have shown that even simple models of plasma physics can be difficult to simulate with high accuracy and precision on NISQ hardware. The reason is that the encoding of the Hamiltonian is relatively dense, requiring on the order of 10-20 gates per time step. Improved gate performance and the application of error suppression and mitigation techniques have improved the performance by a factor of ~ 10 since our initial investigations, but due to the unfavorable scaling of the encoding it remains unclear how to use more than two or three qubits with present capabilities. Hence, larger systems remain out of reach and may ultimately require error-corrected machines to compute.

To perform calculations that make use of larger number of entangled qubits, it is necessary to explore applications that provide a more natural mapping to the hardware topology that keeps the overall gate depth much lower as the circuit width increases. Thus, the simulation of quantum plasma lattice algorithms that communicate through nearest-neighbor interactions and that match the layout of qubits on a grid-based superconducting chip was explored.

Eventually, one might be able to simulate the physics of plasmas through a direct quantum simulation of quantum electrodynamics. However, today's hardware limitations force one to focus on reduced models. In fact, even in the world of classical plasma physics, reduced models are very important for gaining insight into the relevant physical processes. For example, Maxwell's equations require solving for the electric field, \mathbf{E} , and the magnetic field, \mathbf{B} , in terms of the charge density, ϱ , and current density, \mathbf{J} , in three dimensions (3D)

$$\partial_t \varepsilon_0 \mathbf{E} = \nabla \times \mathbf{B} - \mu_0 \mathbf{J} \quad (4.1)$$

$$\nabla \cdot \varepsilon_0 \mathbf{E} = \varrho \quad (4.2)$$

where ε_0 is the vacuum electric permittivity and μ_0 is the vacuum magnetic permeability.

Yet, a full Vlasov-Maxwell simulation would also require evolving the probability distribution function (PDF) for the different particle species in 6D which is very computationally demanding. A reduced model that is much less expensive is to solve the fluid equations

4. Simulating electron-plasma waves on a superconducting quantum chip

for the evolution of moments of the PDF in 3D. At constant temperature, T_s , the fluid moment equations for the number density n_s , and velocity, \mathbf{v}_s , of particle species s are

$$\partial_t n_s + \nabla \cdot (n_s \mathbf{v}_s) = 0 \quad (4.3)$$

$$\partial_t m_s n_s \mathbf{v}_s + \nabla \cdot (m_s n_s \mathbf{v}_s \mathbf{v}_s) + \nabla p_s = e_s n_s E_x \quad (4.4)$$

where the pressure satisfies the ideal gas law, $p_s = n_s T_s$, m_s is the mass, and e_s is the charge of the particles. For 1D, these equations simplify even further. Maxwell's equations reduce to

$$\partial_x E_x = e n_e - e_i n_i \quad (4.5)$$

where $e > 0$ is the magnitude of the electron charge and the isothermal fluid equations reduce to

$$\partial_t n_s + \partial_x n_s v_s = 0 \quad (4.6)$$

$$\partial_t m_e n_e + \partial_x (m_s n_s v_s v_s + p_s) = e_s E_x. \quad (4.7)$$

4.2 Theory

4.2.1 Electron Plasma Wave Model

The first model we will consider is that of electron plasma waves; also known as Langmuir waves. Because the electron mass is much smaller than the ion mass, the ions can be considered to represent a stationary neutralizing background charge density. If we linearize the fluid equations around the fixed background density n_0 and vanishing velocity, then, they reduce to

$$\partial_t n_s + \partial_x n_0 v_s = 0 \quad (4.8)$$

$$\partial_t m_e n_0 v_e + \partial_x n_e T_e = -e E_x. \quad (4.9)$$

The dispersion relation for the linearized equations then yields

$$\omega^2 = \omega_p^2 + k^2 v_{Te}^2 \quad \omega_p^2 = e^2 n_0 / \varepsilon_0 m_e \quad (4.10)$$

where ω is the frequency, k is the wave number, $v_{Te} = \sqrt{T_e/m_e}$ is the thermal speed, and ω_p is the plasma frequency. This represents a system of two partial differential equations (PDEs), that can be discretized using a finite difference approximation, leading to nearest neighbor interactions along the lattice of mesh points. The quantum version of the electron plasma wave model represents a quantum field theory defined on the points of the lattice. In the following, we show that variations of the quantum Heisenberg model can be used to simulate both linear and nonlinear models of electron plasma waves.

4. Simulating electron-plasma waves on a superconducting quantum chip

Large amplitude plasma waves are a key example of resonant nonlinear interactions. Linear interactions imply that the electron velocity is tied to the electric field by the relation

$$\mathbf{v}_e = e\mathbf{E}/im_e\omega \simeq e\mathbf{E}/im_e\omega_p. \quad (4.11)$$

The nonlinearity results from the quivering motion of electrons in the large amplitude oscillations of the electric field. These so-called ponderomotive forces imply that there is a correction to the electron density that takes the form

$$n_e/n_0 = \langle \exp(-m_e \mathbf{v}_e^2/2T_e) \rangle \simeq 1 - \langle m_e |\mathbf{v}_e|^2 / 2T_e \rangle = 1 - \epsilon_0 |\mathbf{E}|^2 / 2p_e. \quad (4.12)$$

Thus, in the reference frame of the wave, the momentum equation becomes

$$i\partial_t E_x = \omega_p E_x - \frac{v_{Te}^2}{2\omega_p} \partial_x^2 E_x - \omega_p \frac{\epsilon_0 |E_x|^2}{4p_e} E_x \quad (4.13)$$

which is the nonlinear Schrödinger equation (NLSE) for $\psi = E_x$. The NLSE is one of the well-known examples of an integrable PDEs that possess soliton solutions and can be solved by the inverse scattering transform.

It is also known that the NLSE is isomorphic to the Heisenberg model, for both classical and quantum variants. In order for a PDE to be integrable, it must have a pair of Lax operators that define translation in time and space for a multi-component wavefunction $\Phi(t, x)$

$$i\partial_t \Phi = \hat{H}\Phi \quad -i\partial_x \Phi = \hat{P}\Phi. \quad (4.14)$$

and commute via the Lax equation

$$\partial_t \hat{P} + \partial_x \hat{H} = i[\hat{P}, \hat{H}]. \quad (4.15)$$

The equations can then be considered to represent a non-Abelian SU(2) gauge theory, where the vector potential that mediates the interactions is a pure gauge field. The equivalence of the NLSE and the Heisenberg model can be demonstrated by showing that they differ by a pure gauge transformation. This was first proven for the classical continuous case by Zakharov and Takhtadzhyan [63] and was extended to the discrete case by Hoffman [64]. In the continuous case, the gauge transformation is

$$S := \mathbf{S} \cdot \boldsymbol{\sigma} = G\sigma_z G^{-1} \quad (4.16)$$

where G is a unitary gauge transformation that is related to the complex wavefunction $\psi = E_x$ via

$$\frac{1}{2}S\partial_x S = G\partial_x G = \text{Re}(\psi)\sigma_x + \text{Im}(\psi)\sigma_y \quad (\partial_x S)^2 = |\psi|^2. \quad (4.17)$$

Thus, the Heisenberg model can also be used to simulate the nonlinear physics of electron plasma waves.



4. Simulating electron-plasma waves on a superconducting quantum chip

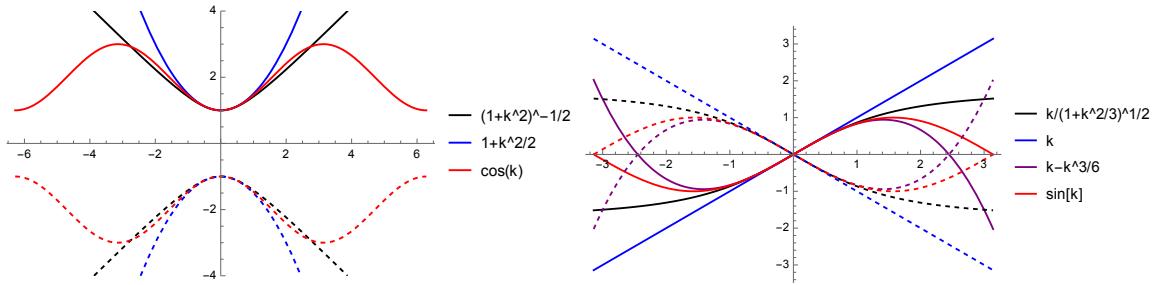


Figure 4.1 – Models for dispersion relations of waves in plasmas (ω vs. k). (Left) Electron plasma waves: exact $(1 + k^2)^{1/2}$ (black), $k^2/2 - 1$ (blue), periodic lattice model $\cos k - 2$ (red) (Right) Ion sound waves: exact $k/(1 + k^2/3)^{1/2}$ (black), k (blue), $k - k^3/6$ (purple), periodic lattice model $\sin k$ (red) Negative frequency branches shown as dashed lines.

Knowledge of the spin vector, S , allows one to recover the complex wavefunction, ψ . In principle, this can be performed using a quantum computer by computing two copies of S , applying the momentum operator to one copy to form $\partial_x S$, and finally forming the product, $S \partial_x S$ as a post-processing step. Similarly, if one can compute the solution to the NLSE, then one can proceed in the opposite direction by computing G and G^{-1} through Hamiltonian simulation of Eq. 4.17 and then forming the product $S = G \sigma_z G^{-1}$ via Eq. 4.16.

4.2.2 Ion Sound Wave Model

The second model we consider is that of ion sound waves; also known as ion acoustic waves. In this case, the electrons respond rapidly to the ions and redistribute their density in order for the plasma to remain quasi-neutral, i.e. so the total charge density is much smaller than the density of the charge carriers themselves, $\varrho/e_i n_i \ll 1$. In this case, electron force balance, $\nabla p_e = e n_e \nabla \phi$, where $\mathbf{E} = -\nabla \phi$, implies that, at constant temperature, T_e , the electrons satisfy the Boltzmann relation $n_e = n_0 \exp(-e_e \phi/T_e)$. Due to quasi-neutrality, summing the electron and ion force balance equations then leads to

$$\partial_t n_i + \nabla \cdot (n_i \mathbf{v}_i) = 0 \quad (4.18)$$

$$\partial_t m_i n_i \mathbf{v}_i + \nabla \cdot (m_i n_i \mathbf{v}_i \mathbf{v}_i) + \nabla (p_i + p_e) = 0. \quad (4.19)$$

For 1D, labeled by coordinate x , this reduces to

$$\partial_t n_i + \partial_x (n_i \mathbf{v}_i) = 0 \quad (4.20)$$

$$\partial_t m_i n_i \mathbf{v}_i + \partial_x (m_i n_i \mathbf{v}_i \mathbf{v}_i) + \partial_x p_i = -e_i n_i \nabla \phi \simeq -\nabla p_e. \quad (4.21)$$

The linear dispersion relation for these waves is

$$\omega^2 = k^2 c_s^2 \quad c_s^2 = (T_i + T_e)/m_i \quad (4.22)$$

where c_s is the ion sound speed and T_s is the temperature of the respective species.

4. Simulating electron-plasma waves on a superconducting quantum chip

If we remove the quasi-neutral assumption, then the difference between the electron and ion density determines the electric potential through the Poisson equation.

$$\nabla^2 \epsilon_0 \phi = e n_e - e_i n_i \quad (4.23)$$

Inserting the expression for the electron pressure then gives the more complete dispersion relation

$$\omega^2 = k^2 c_s^2 / (1 + k^2 \lambda_d^2) \quad \lambda_d^2 = v_{Te}^2 / \omega_p^2 = \epsilon_0 T_e / e^2 n_e \quad (4.24)$$

where λ_d is the Debye length, and we now assume $T_i = 0$ to simplify the rest of the discussion. For wavelengths that are long compared to the Debye length, this has the expansion

$$\omega \simeq k c_s (1 - k^2 \lambda_d^2 / 2 + \dots) \quad (4.25)$$

which represents a third order dispersion of the waves.

The linear physics of the ion sound wave can also be simulated by the Heisenberg model. For the quasi-neutral model, the parameters of the Heisenberg model should be tuned to the critical point, where the effective mass of the excitations vanishes. The fact that the Heisenberg model on a lattice has higher order dispersion, e.g. $\omega = (c/a) \sin(ka) \simeq kc(1 - (ka)^2/6 + \dots)$ can be used to model the dispersive effects of charge separation in Eq. 4.25 by setting $c = c_s$ and $a = \sqrt{3}\lambda_d$. Because the waves in the Heisenberg model only travel in one direction, technically speaking, one must either use two copies of the Heisenberg model or couple two copies together in a manner that generates the desired dispersion relation. For example, one can use the same model as for linear electron plasma waves in the limit that the effective mass tends towards zero.

Coupling the electron plasma and ion sound waves leads to the nonlinear Zakharov equations [65], which couples a linear ion sound model to the nonlinear ponderomotive electron density. This generalizes the ponderomotive potential to include the linear dynamics of the ion sound waves. This system is also integrable and possesses soliton solutions. We are exploring the ways that quantum computers can efficiently solve the Zakharov equations, e.g. by solving two coupled Heisenberg models.

Nonlinear ion sound wave models that include resonant self-interactions are also of interest. If we now transform to a reference frame moving at the sound speed and retain the nonlinear response to the electric force, $-n_i \nabla \phi$, then one obtains the celebrated Korteweg-de Vries (KdV) equation for the ion density

$$2\partial_t n_i + c_s \partial_x n_i^2 / n_0 + c_s \lambda_d^2 \partial_x^3 n_i = 0. \quad (4.26)$$

The KdV equation is another important integrable PDE that possesses soliton solutions and can be solved using the inverse scattering transform. It can also be discretized on the lattice and quantum versions of the model have been formulated. We are exploring quantum algorithms that can solve the quantum lattice KdV model as well.

4. Simulating electron-plasma waves on a superconducting quantum chip

4.2.3 Spin model

The spin model that we use to simulate electron-plasma waves is

$$H = -\frac{J}{4} \sum_j (X_j Y_{j+1} - Y_j X_{j+1}) + \sum_j m_j (-1)^j Z_j. \quad (4.27)$$

Here, J is the strength of the nearest-neighbor spin-spin interaction, and m_j is an onsite field with alternating sign on odd and even sites. For uniform $m_j = m$, the excitation spectrum of this spin model has two branches – positive and negative energy (frequency) branches,

$$\epsilon_k = \pm \sqrt{J^2 \sin^2(ka) + m^2}, \quad (4.28)$$

where a is the lattice spacing. At long wavelength, i.e. small k , we have $\sin(ka) \approx ka$, and the spectrum in the positive branch matches that of electron-plasma waves. The mass gap for the plasmon is m . Therefore, this spin model is a convenient platform to simulate electron-plasma waves. Time evolution of a chain of qubits with this Hamiltonian can be implemented using Trotterization.

If we take the limit that $m \rightarrow 0$, then this also becomes a useful model for linear ion sound waves. As discussed earlier, with the approximation $\sin(ka) \simeq ka[1 - (ka)^2/3! + \dots]$, one can make the dispersion relation approximately match the ion sound wave dispersion relation, $kc/(1 + (k\lambda_d)^2)^{1/2} \simeq kc[1 - (k\lambda_d)^2/2 + \dots]$, if one choose the lattice spacing to match the Debye length, $a = \lambda_d/\sqrt{3}$.

4.3 Electron-plasma wave simulations with the spin model

We perform three kinds of experiments that simulate evolution with the above spin model. In one experiment, we demonstrate that the spin model's spectrum matches the electron-plasma wave's spectrum. In the next two experiments, we simulate the propagation of electron-plasma waves that have the potential to be useful in specific applications. Our experiments were done on Rigetti's Ankaa-9Q-3 device.

4.3.1 Experimental scheme to measure the spectrum

To verify that the spectrum for H is given by Eq. (4.28), we perform a many-site Ramsey-type experiment. First, we prepare one qubit in $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, and all the other qubits in $|1\rangle$. This prepares a superposition of the vacuum state, $|\psi_{\text{vac}}\rangle$, and the state with one wavepacket localized on one site, $|\psi_{1\text{-ex}}\rangle$,

$$|\psi_0\rangle = \frac{|\psi_{\text{vac}}\rangle + |\psi_{1\text{-ex}}\rangle}{\sqrt{2}}. \quad (4.29)$$



4. Simulating electron-plasma waves on a superconducting quantum chip

We then implement Trotter evolution with H . The logical circuit which implements this experiment is shown in Fig. 4.2(a).

The localized excitation is a sum over all the single-excitation eigenmodes $|k\rangle$ of the Hamiltonian,

$$|\psi_{1\text{-ex}}(t=0)\rangle = |\cdots 00100\cdots\rangle = \sum_k c_k |k\rangle. \quad (4.30)$$

Each eigenmode accrues a time-dependent phase that is proportional to its eigenvalue,

$$|\psi_{1\text{-ex}}(t)\rangle = \sum_k c_k e^{-i\epsilon_k t} |k\rangle. \quad (4.31)$$

The vacuum state doesn't evolve, since it is a zero-energy eigenstate of the Hamiltonian. We are interested in obtaining the eigenvalues ϵ_k . To that end, we measure $\langle X_j \rangle + i \langle Y_j \rangle$. Let us denote the real-space wavefunction amplitudes of $|k\rangle$ as Λ_{kj} . It can be shown that $\langle k|(X_j + iY_j)|\psi_{\text{vac}}\rangle = \Lambda_{kj}$. Moreover, we have the relations $\langle k|(X_j + iY_j)|k'\rangle = \langle \psi_{\text{vac}}|(X_j + iY_j)|k\rangle = \langle \psi_{\text{vac}}|(X_j + iY_j)|\psi_{\text{vac}}\rangle = 0$. Using these, we find that

$$\langle X_j + iY_j \rangle (t) = \sum_k c_k^* e^{i\epsilon_k t} \Lambda_{kj} \quad (4.32)$$

Using the orthogonality of Λ_{kj} , we arrive at

$$e^{i\epsilon_p t} = \frac{\sum_j \Lambda_{pj}^* \langle X_j + iY_j \rangle (t)}{c_p^*} \quad (4.33)$$

from which we can estimate ϵ_p .

While it should be noted that the above technique of verifying ϵ_p from a many-body Ramsey-type experiment requires to classically calculate the Λ_{kj} , the Λ_{kj} are easy to calculate here since the spin model is exactly solvable. For general spin models, more advanced techniques such as quantum phase estimation may be required.

4.3.2 Experimental scheme to simulate injection of plasma waves

Doppler backscattering is a type of reflectometry technique that is quite useful for obtaining sensitive measurements of plasma density. In this case, a wave-packet is reflected by the steep rise in electron density that occurs at the edge of a confined plasma. Thus, we performed an experiment aimed at modeling similar processes using a quantum computer. Specifically, we simulated injection of wave packets into a plasma, by time-evolving the qubits under the spin Hamiltonian with a sharp domain in the on-site field m_i .

To perform this experiment, we initially prepared an entangled state of the qubits,

$$|\psi\rangle = (\alpha|01\rangle + \beta|10\rangle) \otimes |11\cdots\rangle \quad (4.34)$$

4. Simulating electron-plasma waves on a superconducting quantum chip

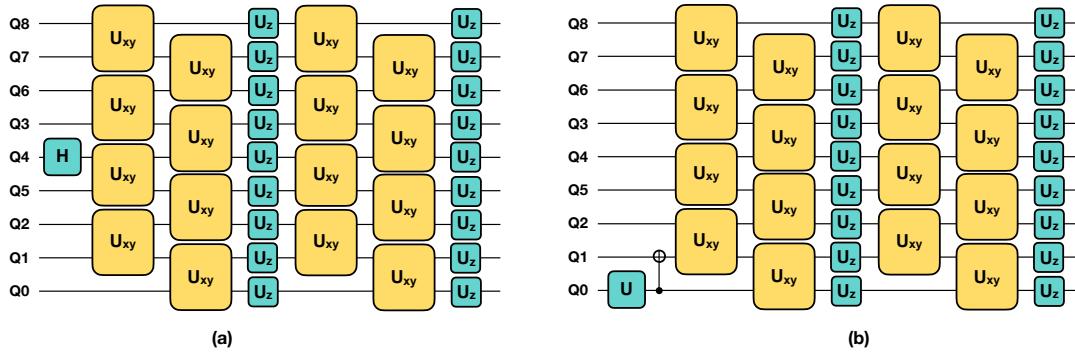


Figure 4.2 – (a) Logical circuit used to measure the dispersion of excitations in the spin model that simulates electron-plasma waves [Eq. (4.27)]. (b) Logical circuit used to simulate the propagation of waves through an inhomogeneous plasma or a plasma with a sharp jump in electron density. In these circuits, $U_{xy}^{j,j+1} = \exp\left(i\frac{J\delta t}{4}(X_j Y_{j+1} - Y_j X_{j+1})\right)$ and $U_z = \exp(-im_j\delta t(-1)^j Z_j)$, where δt is the Trotter time step, J is the strength of the XY interaction in the spin model, and m_j is the onsite field. The one-qubit gate U in (b) is such that $U|0\rangle = \alpha|0\rangle + \beta|1\rangle$ (see text for definition of notations).

where α and β are complex coefficients normalized to 1. We then implement time evolution under H with $m_j = \Theta(j - N/2)$, where Θ is the Heaviside theta function. We measure $\langle Z_j \rangle$ at the end. The logical circuit is shown in Fig. 4.2(b).

This experiment simulates a wave packet, initially localized to two sites, and propagating through the lattice. The local wave packet has a broad distribution in momentum space, due to the Heisenberg uncertainty principle. A more general W state, $|\psi\rangle = \sum_j \alpha_j |\cdots 11\rangle |0\rangle_j |11\cdots\rangle$ which has a broader real-space envelope, would be more sharply peaked in momentum. The wave packet's density is $(1 - \langle Z_j \rangle)/2$. In the ideal limit of a wave packet with a precise momentum, it gets injected into the plasma only if its mean energy $J \sin(ka)$ is larger than the mass gap m .

4.3.3 Propagation in an inhomogeneous medium

Waves are used as important actuators in plasma experiments because they can be used to control the injection of energy and momentum. Simulations of wave propagation in an inhomogeneous plasma are needed to predict where the waves deposit their energy and how the controller will perform in practice. A key issue is that waves tend to scatter when they travel through an inhomogeneous medium, and, for example, this is true for electron plasma waves that travel through a plasma with an inhomogeneous electron density. Thus, we designed an experiment to explore how electron plasma waves scatter off of a peaked plasma density profile. Specifically, we simulate propagation of waves through an inhomogeneous plasma medium, by time-evolving the qubits under the spin Hamiltonian with a peaked profile for the m_i .

This experiment is similar to the prior experiment, except that m_j has a peaked profile

in space. We choose $m_j = m_0 \exp(-(j - N/2)^2/\delta^2)$.

4.4 Running an experiment on Ankaa-9Q-3

Rigetti's Ankaa-9Q-3 QPU offers digital one-qubit $\text{RX}(\pi)$, $\pm\text{RX}(\pi/2)$, and $\text{RZ}(\theta)$ gates, and a digital ISWAP gate on a square lattice¹. The analog part of the scheme we employ for electron-plasma wave simulation is the entangling FSIM gate. The FSIM describes, to an excellent approximation, the native Hamiltonian of the device when the tunable coupler between two qubits is activated. This interaction is higher fidelity than ISWAP; we have previously achieved a two-qubit gate fidelity of 99.4% for an ISWAP-like FSIM [31]. Combined with the digital one-qubit gates, FSIM provides a universal gate set to which any quantum algorithm can be compiled.

Digital one-qubit gates are generally well-understood and simple to optimize within the coherence limits of the qubits without optimal control methods [66]. Arbitrary single-qubit unitary operations can then be realized accurately by using three phased microwave pulses [67] known as the "PMW3" decomposition. To maximize the performance of the electron-plasma wave simulation, we analyze the expressiveness of candidate gates in achieving the required two-qubit interaction unitaries [56]. As in previous work [2] we decided to target a square-root-of-ISWAP ($\sqrt{\text{ISWAP}}$) gate. We modify the digital ISWAP gate using Rigetti's pulse-level control language, Quil-T², to obtain the a $\sqrt{\text{ISWAP}}$ -like FSIM. We learn each FSIM unitary using process tomography and a variant of cross-entropy benchmarking [68, Supplemental].

A $\sqrt{\text{ISWAP}}$ -like FSIM gate and one-qubit gates can be used to construct any two-qubit operation for quantum computation. We use a numerical compilation technique to express the target logical circuits in terms of the characterized FSIM gates and arbitrary one-qubit gates. For example, we express the U_{xy} gate in Fig. 4.2 approximately in terms of two FSIM gates and single-qubit gates, because this confers a higher fidelity than an exact decomposition to three FSIM gates. Our one-qubit gates are implemented using the PMW3 decomposition, the details of which are provided in Appendix A.1. This results in the hardware-native experimental circuit illustrated in Fig. 4.3.

Rigetti's parametric compilation and execution support can be leveraged to express the parameters of this circuit in terms of only digital phases in the one-qubit gate decompositions. This supports us to run various error mitigation strategies over the hardware-native circuit and to obtain the highest-quality algorithm-level result. We make use of our FPGA-accelerated randomized readout [69] as part of a broader strategy to tailor noise on the device and make error mitigation more effective. Tailoring is achieved using a variant of twirling [70]. We then implement Clifford Data Regression similar to [71] to rescale expectation values. We leverage knowledge of the electron-plasma wave

¹For current device specifications, refer to <https://qcs.rigetti.com/qpus>.

²See the Rigetti Computing article "Gain deeper control of Rigetti quantum processors with Quil-T".

4. Simulating electron-plasma waves on a superconducting quantum chip

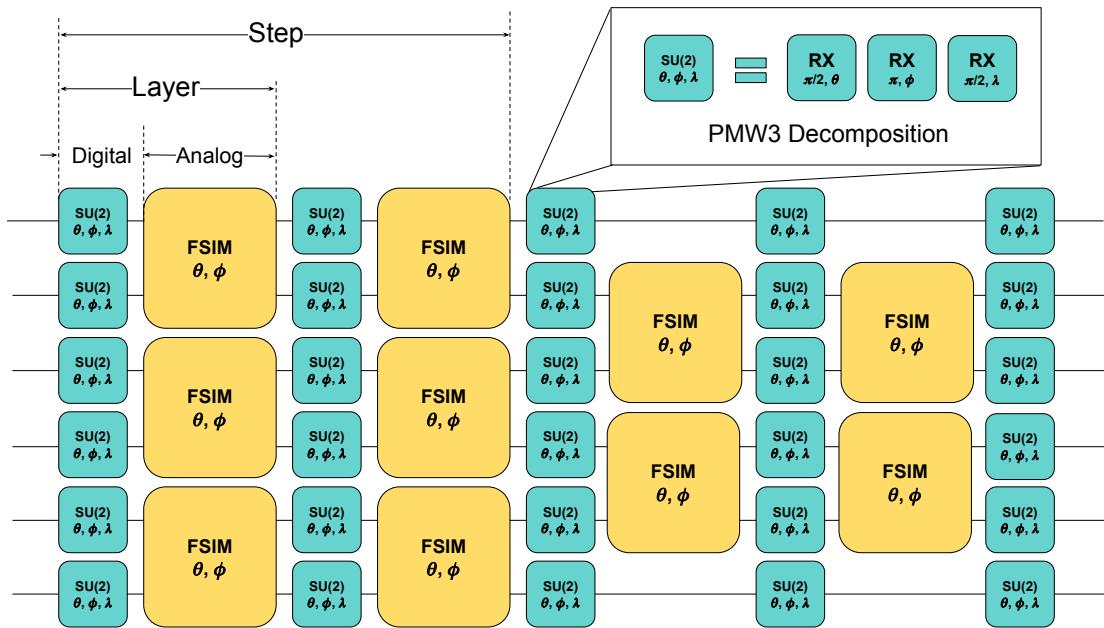


Figure 4.3 – Hardware-native circuit that performs the experiments in Section 4.3. The Trotter time steps are broken down into layers, and then into digital 1Q and analog 2Q gates. Each 1Q gate is implemented via a PMW3 decomposition to digital RX and RZ gates. Each FSIM gate is learned using process tomography and cross-entropy benchmarking.

problem, where evolution with the spin Hamiltonian is number-conserving, to fine-tune the rescaling coefficients similar to the post-selection approach of [72].

4.5 Experimental results

Experiments were performed on the Ankaa-9Q-3 QPU with an error per layered gate (EPLG) [73] of 1.7% (see Table A.1 for other performance data). Larger experiments at the 30-qubit scale were initially envisioned, however resolving the plasma dynamics of interest requires a relatively high circuit depth. With current error rates, it is not feasible to resolve observables for electron-plasma wave simulation at a scale of 30 qubits (see Appendix A.2). The experiment was thus scaled to 9 qubits where high-depth circuits can be viably executed and measured. It's estimated that a further reduction in EPLG to 0.5% is necessary to scale the experiment to 30 qubits. Fortunately, the control, compilation and mitigation techniques utilized in our experiments are scalable, and should enable simulations of this size with commensurate reductions in error.

4. Simulating electron-plasma waves on a superconducting quantum chip

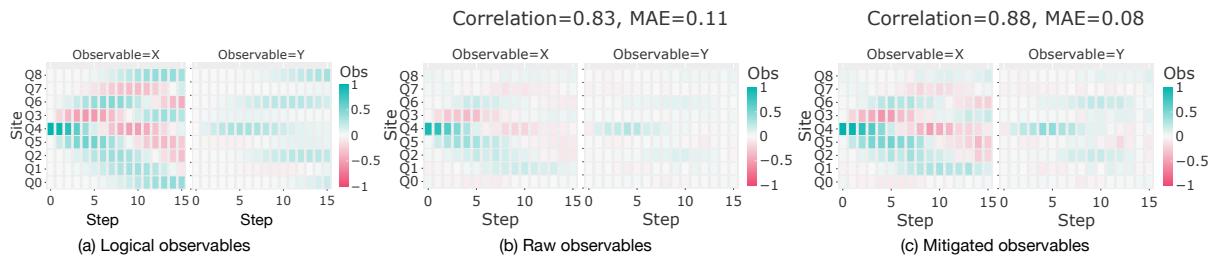


Figure 4.4 – The time evolution of $\langle X \rangle$ and $\langle Y \rangle$ is shown in this many-site Ramsey-type experiment. The time evolution of the complex phase of $\langle X + iY \rangle$ reveals the energies of all the electron-plasma wave eigenmodes. The subplots show the evolution without noise (left), the raw estimated observables (middle), and the error-mitigated results (right).

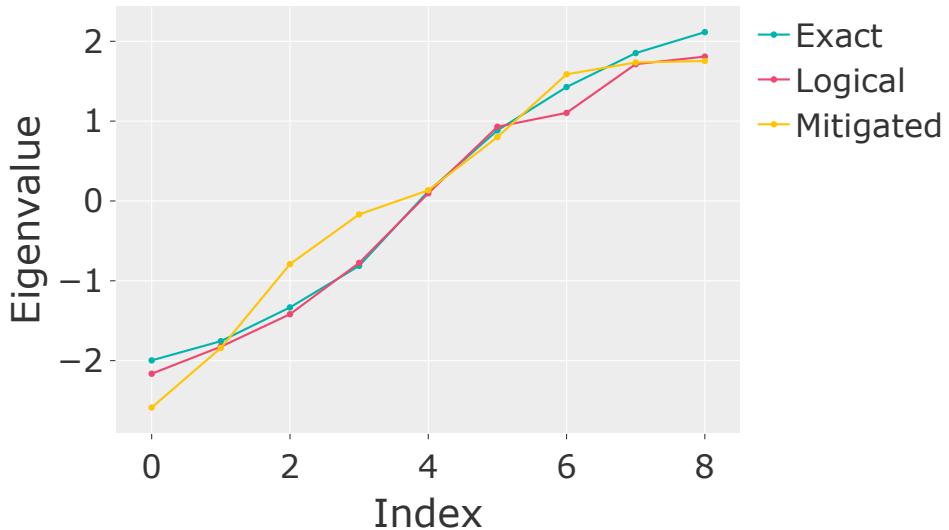


Figure 4.5 – The dispersion relation for the electron-plasma wave, simulated with a 1D spin model. The teal curve plots the exact eigenvalues of the 1D spin model – equivalently the energies of the electron-plasma wave eigenmodes, the magenta curve the result obtained assuming a noiseless experiment, and the yellow curve the result from our error-mitigated observables.

4.5.1 Dispersion

Fig. 4.4 plots the evolution of $\langle X_j \rangle$ and $\langle Y_j \rangle$ with time, in an experiment with $J = 1$, $m = 0.25$, and Trotter time step $\delta t = 0.4$. We executed circuits up to 16 time steps. The initial state is a superposition of the vacuum state and a wave packet localized on one site, simulated by preparing one qubit in the $|+\rangle$ state. Fig. 4.4 shows the wave packet spreading with time, and the experimental measurement of the same (Fig. 4.4c) agrees with the predicted behavior (Fig. 4.4a) after error mitigation, with a Pearson correlation $C = 0.88$ and mean absolute error of 0.08. This evolution encodes all the eigen energies

4. Simulating electron-plasma waves on a superconducting quantum chip

of the wave packet, which can be extracted from the data. Fig. 4.5 plots the extracted eigenvalues, and shows good agreement between the results from mitigated data and the predicted result.

4.5.2 Scattering of Waves in an Inhomogeneous Plasma Medium

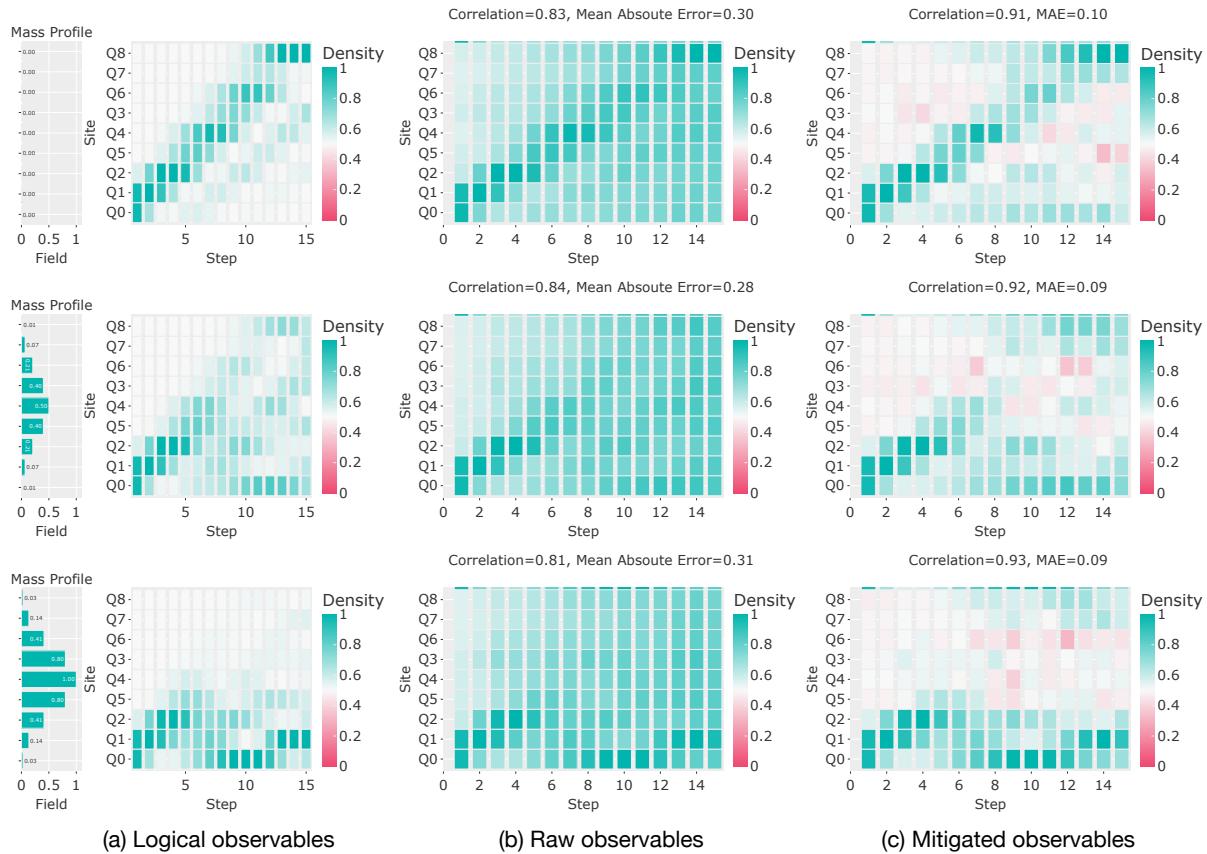


Figure 4.6 – This experiment models an electron plasma wave packet propagating in a inhomogeneous plasma. For each row, the plasma density, i.e. the effective mass profile, is shown on the left. The subplots from top to bottom show the result with an increasing peak density, i.e. barrier height. The simulated circuits show the exact evolution without noise (left), the raw observables estimated from experimental results in the presence of noise (middle), and the mitigated results where the observables are estimated after rescaling (right).

Fig. 4.6 shows the evolution of a wave packet that is initially two sites wide, $|\psi_0\rangle = (\alpha|01\rangle + \beta|10\rangle) \otimes |1\rangle^{\otimes(N-2)}$. The wave packet then propagates through a plasma with a Gaussian density profile that leads to the spatially dependent effective mass gap m_i shown to the left of each row. The Trotter time step is $\delta t = 0.8$. The height of the effective barrier increases from top to bottom. The error-mitigated experimental data captures the dynamics reasonably well, with a Pearson correlation of $\gtrsim 0.9$ with the

4. Simulating electron-plasma waves on a superconducting quantum chip

predicted values. The reflection of the wave packet off the barrier increases as the barrier height increases.

4.5.3 Reflection of Waves from the Plasma Edge

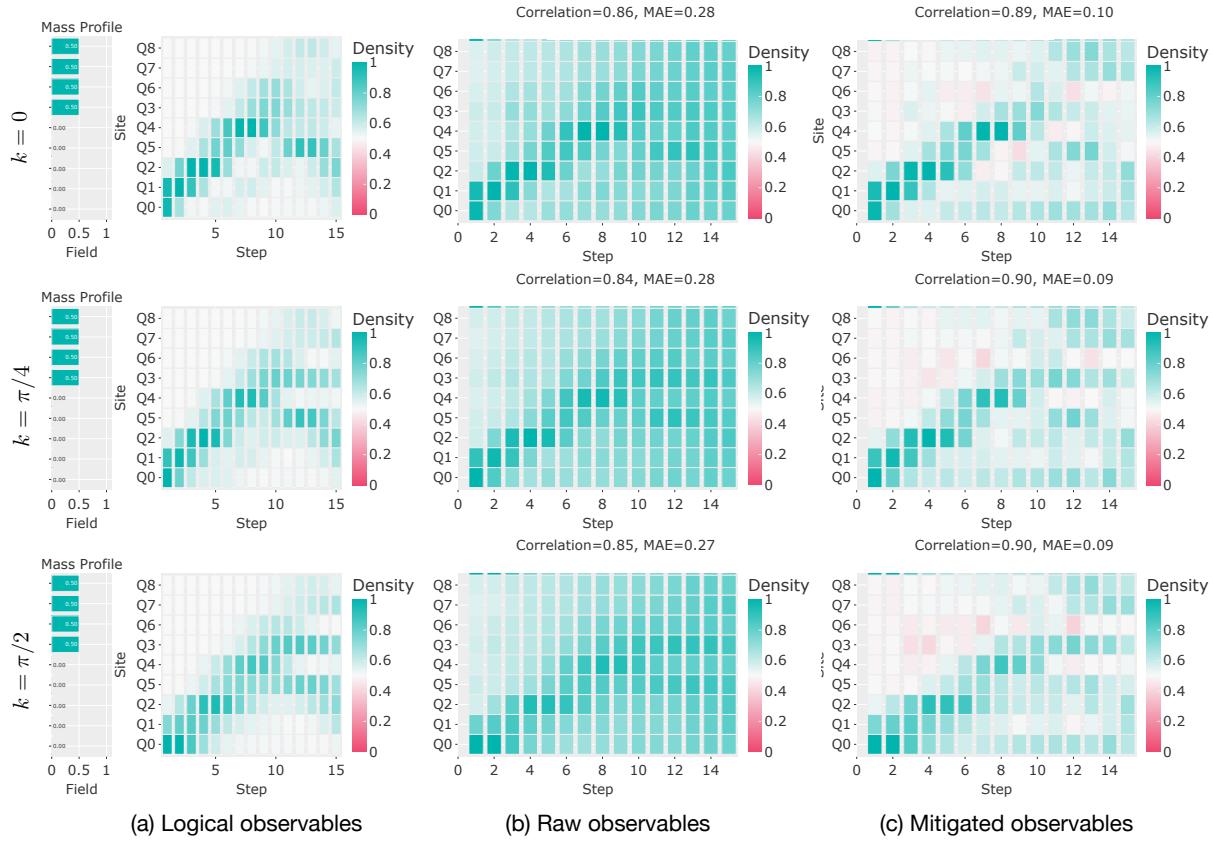


Figure 4.7 – This experiment models an electron plasma wave packet encountering a sharp jump in density, which naturally at the edge of a confined plasma. In each row, the plasma density profile, i.e. the effective mass barrier, is shown on the left. The subplots from top to bottom show the result with different initial values of wavenumber, k , which are set by the initial phase of the wave-packet. The simulated circuits show the exact evolution without noise (left), the raw observables estimated from experimental results in the presence of noise (middle), and the mitigated results where the observables are estimated after rescaling (right).

In this case, we simulated the reflection of an electron plasma wave packet that encounters a sharp jump in plasma density. Fig. 4.7 shows the evolution of a wave packet that is initially two sites wide, $|\psi_0\rangle = (\alpha|01\rangle + \beta|10\rangle) \otimes |1\rangle^{\otimes(N-2)}$, and launched into a homogeneous plasma medium with a step-like profile in plasma density, i.e. the effective mass ($m_i = \frac{1}{2}\Theta(i-5)$). The Trotter time step is $\delta t = 0.8$. In momentum (k) space, the wave packet is distributed over a range of k , with the mean momentum determined by the relative phase between α and β , albeit the k -dependence being limited on these 9Q experiments. The amount of reflected wave depends on the mean energy of the

4. Simulating electron-plasma waves on a superconducting quantum chip

wave packet, which depends on k . The error-mitigated experimental data again captures the dynamics reasonably well, with a Pearson correlation of $\gtrsim 0.89$ with the predicted values.

4.6 Conclusion

By incorporating advanced control, a digital-analog gate-scheme with numerical compilation and a suite of error mitigation techniques, the dynamics of electron-plasma waves were simulated up to 15 time steps on 9 qubits. The dynamics of electron plasma wave-propagation in an inhomogeneous plasma medium, and reflection and transmission at the edge of a plasma, represented by a barrier in effective mass, were successfully recreated on a superconducting quantum computer. While the scale of the experiment is well within the realm of classical simulation, the techniques applied in the experiment are scalable to the level of hundreds of qubits and investigations of potential quantum advantage. Moreover, the exploration here illustrates a viable path towards more general simulation capabilities. Both parts of the analog-digital case scheme used in this experiment have clear paths to improvement. The 1Q gate speed and accuracy can be improved through better control, allowing for reduced errors in the digital portion. The analog step is currently performed on layers of disjoint pairs of qubits. These layers could be made more expressive by allowing multi-qubit interactions, allowing the simulation of more diverse Hamiltonians with richer interactions and connection topology. The numerical approach to expression has proven flexible and accurate, but careful attention must be paid to performance as the parameter space grows - the field of GPU-accelerated linear algebra frameworks may provide promising tools. Finally, coherence must be improved to allow longer and bigger simulations to take place, even as the speed of interaction is accelerated.

Appendix A

Supplemental Material

A.1 The PMW3 decomposition

Rigetti Ankaa-class devices can realize $\text{RX}(\pi)$ and $\pm\text{RX}(\pi/2)$ gates, defined as

$$\begin{aligned}\text{RX}(\pi) &= \begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}, \\ \text{RX}(\pi/2) &= \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-i}{\sqrt{2}} \\ \frac{-i}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}.\end{aligned}$$

In addition, Ankaa-class devices can realize a parametric $\text{RZ}(\theta)$ gate, also known as $\text{PHASE}(\theta)$, using local updates of in-sequence phases that consume zero runtime and introduce negligible error [74],

$$\text{RZ}(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{+i\theta/2} \end{pmatrix}.$$

Using these gates, more complex rotations can be performed in the 1Q space. Of interest here are the $\text{PHASEDSX}(\phi)$ and $\text{PHASEDX}(\phi)$ gates,

$$\begin{aligned}\text{PHASEDSX}(\phi) &\equiv \text{RZ}(-\phi)\text{RX}(\pi/2)\text{RZ}(\phi) = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-ie^{i\phi}}{\sqrt{2}} \\ \frac{-ie^{-i\phi}}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}, \\ \text{PHASEDX}(\phi) &\equiv \text{RZ}(-\phi)\text{RX}(\pi)\text{RZ}(\phi) = \begin{pmatrix} 0 & -ie^{i\phi} \\ -ie^{-i\phi} & 0 \end{pmatrix}.\end{aligned}$$

Using a combination of two $\text{PHASEDSX}(\phi)$ gates and one $\text{PHASEDX}(\phi)$ gate, it's possible to perform any $\text{SU}(2)$ rotation. This is called a PMW3 decomposition, since it uses three phased microwave pulses. Moreover, while virtual-Z gate schemes require commuting Z-rotations through 2-qubit gates, the PMW3 scheme does not. Although the phases of the microwave pulses can be controlled digitally, all composite rotations are physical.

$$\begin{aligned}
 PMW3(\theta, \phi, \omega) &= PHASEDSX(\theta)PHASEDX(\phi)PHASEDSX(\omega) \\
 &= \begin{pmatrix} -e^{i(\theta-\omega)/2} \cos \frac{2\phi-\theta-\omega}{2} & e^{i(\theta+\omega)/2} \sin \frac{2\phi-\theta-\omega}{2} \\ -e^{-i(\theta+\omega)/2} \sin \frac{2\phi-\theta-\omega}{2} & -e^{i(\omega-\theta)/2} \cos \frac{2\phi-\theta-\omega}{2} \end{pmatrix}. \tag{A.1}
 \end{aligned}$$

A.2 Error budget

All NISQ algorithms work within an error budget. That is, given the error rates of the gates and the sampling rate of the device, there is a floor on the smallest coherent signal we can measure. For example, a state-of-the-art superconducting quantum computer may offer 2-qubit gates with 99.5% fidelity. Given a circuit containing 1000 2-qubit gates, we may therefore expect the circuit fidelity to be $0.995^{1000} = 0.00665 = 0.665\%$. In other words, to have a greater than 50-percent chance of observing an error-free bitstring, we would require a little over 100 samples. The sample requirement increases exponentially with the size of the circuit, so for a circuit with 2000 gates, we require about 15,000 samples and a circuit with 4,000 gates would require 350 million samples. Thus, the combination of the error rate and the sample rate determine the size of the circuits which can be practically used in experiments.

The above calculation is very simplistic but provides a reasonable intuition. One important simplification is the assumption that noise is Markovian and depolarizing. These assumptions are true only to first-order in contemporary quantum devices, and the above calculation thus represents a best-case scenario. Moreover, in many cases we do not require a full bitstring but rather are attempting to estimate low-weight observables. This greatly reduces the requirements on fidelity. On the other hand, we are usually interested in estimating the observable to some analog precision and thus require more than a single correct bitstring. The uncertainty of the estimated observable will scale with the $1/\sqrt{N}$ where N is the number of samples. If the observable value is suppressed by errors, to say 10% of its error-free value, we will require $1/0.1^2 = 100$ times the number of samples to achieve the same precision in our estimate. Thus the number of required samples in the presence of noise can quickly become very large.

With the basic requirements of measuring noisy states, we can discuss more concretely the error budget of our problem. In order to perform an interesting simulation, we need a large number of time steps. Each time step requires some number of gates implemented and must be executed in sequence. The number of time steps thus sets the depth of the circuit (M). We also wish to observe the dynamics over large system sizes which requires a number of qubits (N). The total number of gates in the circuit thus scales as $M \times N$.

To construct the circuit, we pick a 1D linear chain topology and apply layers of 1Q and 2Q gates. To perform 1 time step, we require 4 cycles of 1-qubit + 2-qubit gates. For a relatively modest time evolution of 15 steps we require a circuit depth of 60 cycles,

A. Supplemental Material

which is similar to the depth of the circuits used in utility-scale experiments [75]. In order to resolve observables with sufficient precision and a practical number of samples (100,000), a circuit fidelity of around 1% is required. Based on our analysis shown in Fig. A.1, we estimate that an EPLG around 1.75% is required to support an experiment using $N = 9$ qubits. To support $N = 30$ qubits, an EPLG around 0.5% would be required.

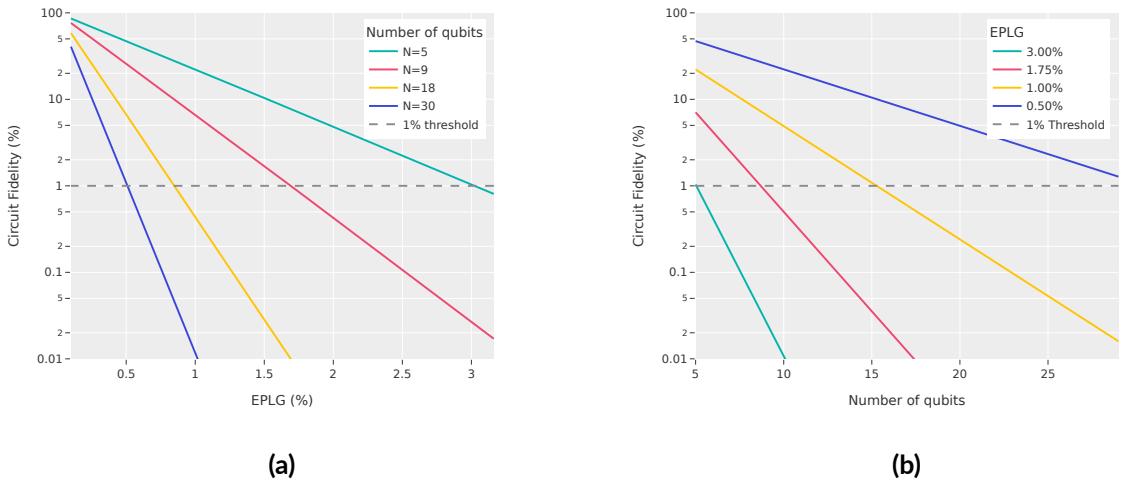


Figure A.1 – Scaling of circuit fidelity with (a) the error per layer gate (EPLG) for different numbers of qubits, and with (b) number of qubits for different EPLG. Experiments using 100,000 samples on $N = 9$ qubits lead to a target circuit fidelity $\sim 1\%$ and a EPLG $\sim 1.75\%$.

A.2.1 Device design and architecture

The Rigetti Ankaa-9Q-3 QPU contains 9 transmon qubits connected via floating tunable couplers [31] arranged in square grid. The tunable couplers enable control of qubit-qubit coupling between neighboring qubits by threading external magnetic flux through the SQUID loops. This allows the coupling to be turned off when qubits are idling or undergoing single-qubit rotations. The coupling can be quickly turned on by applying a base-band flux pulse to the coupler when actuating a two-qubit gate.

Entangling gates between neighboring qubits are activated using two schemes: parametric resonance [76] and fast DC [77]. The parametric resonance gate is activated by bringing the two qubits into resonance with a modulated flux pulse applied to the higher-frequency qubit. This pulse brings it into resonance with its neighbor through an XX+YY type interaction, which is a generator of the iSWAP-family of gates. This type of gate is called the FSIM gate. A base-band flux pulse is simultaneously applied to the coupler to enhance the coupling between the qubits. The fast DC gate works in fundamentally the same way, but the modulated flux pulse is replaced with a base-band flux pulse to bring the qubits into resonance.

A. Supplemental Material

Metric	Value
T1	$28.6\mu s$
T2	$19.7\mu s$
ISWAP Infidelity	1.33%
ISWAP Duration	104ns
RX Infidelity	0.10%
RX Duration	60ns

Table A.1 – Summary metrics of the Ankaa-9Q-3 QPU. All values are median values over the entire processor.

Device architecture

The device employed in our experiments is the Rigetti Ankaa-9Q-3 QPU. The device consists of 9 superconducting tunable transmon qubits connected by tunable couplers in a grid architecture. ISWAP gates are calibrated on the device.

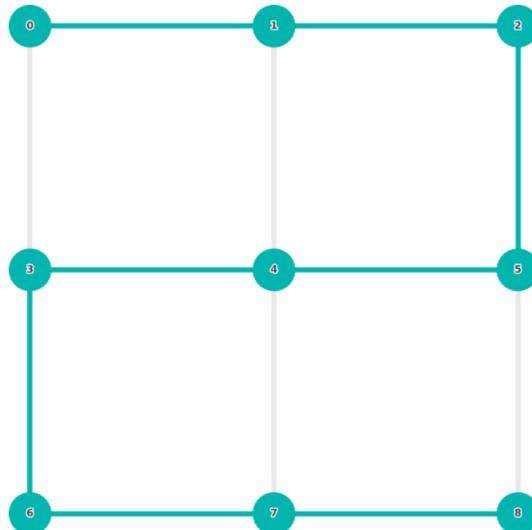


Figure A.2 – The Ankaa-9Q-3 device, annotated with the linear chain lattice used for experiments. The qubits and edges used in this experiment are highlighted in teal.

Device characteristics

The characteristics of the Ankaa-9Q-3 device at the time of testing are shown in Table A.1, as reported on Rigetti Quantum Cloud Services and by analysis of production calibration data. For the experimental configuration of circuits executed in this report, the Error Per Layered Gate (EPLG) was measured to be around 1.7%; sufficient to support the experiments as per the error budget specified above.

References

- [1] Y. Cho. "Direct pulse-level compilation of arbitrary quantum logic gates on superconducting qutrits". *Physical Review Applied* 22.3 (2024). DOI: [10.1103/PhysRevApplied.22.034066](https://doi.org/10.1103/PhysRevApplied.22.034066).
- [2] Y. Shi, B. Evert, A. F. Brown, et al. "Simulating nonlinear optical processes on a superconducting quantum device". en. *Journal of Plasma Physics* 90.6 (2024), p. 805900602. DOI: [10.1017/S0022377824001326](https://doi.org/10.1017/S0022377824001326). URL: <https://www.cambridge.org/core/journals/journal-of-plasma-physics/article/simulating-nonlinear-optical-processes-on-a-superconducting-quantum-device/AE97AB85141FCF53189399402A9A8C3B> (visited on 11/19/2024).
- [3] T. Schenkel, A. Baczewski, M. Boshier, et al. *Fusion Energy Sciences Roundtable on Quantum Information Science*. 2018.
- [4] P. Bonoli and L. C. McInnes. *Report of the Workshop on Integrated Simulations for Magnetic Fusion Energy Sciences*. DOE Office of Science Workshop Report. 2015.
- [5] L. Lao. *Extreme-scale computing with emphasis on high-fidelity physics leading to reduced models for plasma simulation and control technology*. TEC whitepaper. 2017.
- [6] A. Engel, G. Smith, and S. E. Parker. "Quantum Algorithm for the Vlasov Equation". *Phys. Rev. A* 100 (2019), p. 062315.
- [7] J. B. Parker and I. Joseph. "Quantum phase estimation for a class of generalized eigenvalue problems". *arXiv:2002.08497 [quant-ph]* (2020). URL: <https://arxiv.org/abs/2002.08497>.
- [8] I. Joseph. "Koopman-von Neumann Approach to Quantum Simulation of Nonlinear Classical Dynamics". *arXiv:2003.09980 [quant-ph]* (2020). URL: <https://arxiv.org/abs/2003.09980>.
- [9] I. Y. Dodin and E. A. Startsev. "On applications of quantum computing to plasma simulations". *arXiv:2005.14369 [quant-ph]* (2020). URL: <https://arxiv.org/abs/2005.14369>.
- [10] Y. Shi, A. R. Castelli, I. Joseph, et al. "Quantum computation of three-wave interactions with engineered cubic couplings". *arXiv preprint arXiv:2004.06885* (2020). URL: <https://arxiv.org/abs/2004.06885>.
- [11] X. Wu, S. L. Tomarken, N. A. Petersson, et al. "High-fidelity software-defined quantum logic on a superconducting qudit". *arXiv:2005.13165 [quant-ph]* (2020). URL: <https://arxiv.org/abs/2005.13165>.

References

- [12] J. S. Otterbach, R. Manenti, N. Alidoust, et al. “Unsupervised Machine Learning on a Hybrid Quantum Computer” (2017). arXiv:1712.05771.
- [13] I. M. Georgescu, S. Ashhab, and F. Nori. “Quantum simulation”. *Rev. Mod. Phys.* 86 (1 2014), pp. 153–185. DOI: [10.1103/RevModPhys.86.153](https://doi.org/10.1103/RevModPhys.86.153). URL: <https://link.aps.org/doi/10.1103/RevModPhys.86.153>.
- [14] R. C. Davidson. *Methods in Nonlinear Plasma Theory*. Academic Press, 1972.
- [15] S. Baalrud, N. Ferraro, L. Garrison, et al. *A Community Plan for Fusion Energy and Discovery Plasma Sciences: Report of the 2019–2020 American Physical Society Division of Plasma Physics Community Planning Process*. 2020. URL: <https://sites.google.com/ppl.gov/dpp-cpp>.
- [16] D. R. Nicholson. *Introduction to Plasma Theory*. John Wiley & Sons, 1983.
- [17] R. P. Feynman. *Int. J. Theor. Phys.* 21 (1982), p. 467.
- [18] S. Lloyd. “Universal Quantum Simulators”. *Science* 273.5278 (1996), pp. 1073–1078. URL: <http://science.sciencemag.org/content/273/5278/1073>.
- [19] D. W. Berry, A. M. Childs, and R. Kothari. 2015.
- [20] D. W. Berry. 2017.
- [21] P. C. S. Costa, S. Jordan, and A. Ostrander. 2019.
- [22] E. Farhi, J. Goldstone, D. Gosset, et al. “Quantum Adiabatic Algorithms, Small Gaps, and Different Paths”. *Quantum Info. Comput.* 11.3 (2011), pp. 181–214. URL: <http://dl.acm.org/citation.cfm?id=2011395.2011396>.
- [23] E. Farhi and A. W. Harrow. “Quantum Supremacy through the Quantum Approximate Optimization Algorithm”. *arXiv:1602.07674* (2016). URL: <https://arxiv.org/abs/1602.07674>.
- [24] A. Montanaro. *Proc. R. Soc. A* 471 (2015), p. 0301.
- [25] D. A. Lidar and O. Biham. “Simulating Ising spin glasses on a quantum computer”. *Physical Review E* 56.3 (1997), pp. 3661–3681. DOI: [10.1103/PhysRevE.56.3661](https://doi.org/10.1103/PhysRevE.56.3661). URL: <https://link.aps.org/doi/10.1103/PhysRevE.56.3661>.
- [26] A. Zlokapa, S. Boixo, and D. Lidar. *Boundaries of quantum supremacy via random circuit sampling*. 2020.
- [27] M. Kjaergaard, M. E. Schwartz, J. Braumüller, et al. “Superconducting Qubits: Current State of Play”. *Annual Review of Condensed Matter Physics* 11.1 (2020), pp. 369–395. DOI: [10.1146/annurev-conmatphys-031119-050605](https://doi.org/10.1146/annurev-conmatphys-031119-050605). URL: <https://doi.org/10.1146/annurev-conmatphys-031119-050605>.
- [28] S. A. Caldwell, N. Didier, C. A. Ryan, E. A. Sete, A. Hudson, P. Karalekas, R. Manenti, M. P. da Silva, R. Sinclair, E. Acala, et al. “Parametrically Activated Entangling Gates Using Transmon Qubits”. *Phys. Rev. Applied* 10 (3 2018), p. 034050. DOI: [10.1103/PhysRevApplied.10.034050](https://doi.org/10.1103/PhysRevApplied.10.034050).

References

[29] M. Reagor, C. B. Osborn, N. Tezak, et al. "Demonstration of universal parametric entangling gates on a multi-qubit lattice". *Science Advances* 4 (2018), eaao3603. DOI: [10.1126/sciadv.aao3603](https://doi.org/10.1126/sciadv.aao3603).

[30] N. Didier, E. A. Sete, M. P. da Silva, and C. Rigetti. "Analytical modeling of parametrically modulated transmon qubits". *Phys. Rev. A* 97 (2 2018), p. 022330. DOI: [10.1103/PhysRevA.97.022330](https://doi.org/10.1103/PhysRevA.97.022330).

[31] E. A. Sete, A. Q. Chen, R. Manenti, S. Kulshreshtha, and S. Poletto. "Floating tunable coupler for scalable quantum computing architectures". *Physical Review Applied* 15.6 (2021). arXiv:2103.07030 [quant-ph], p. 064063. DOI: [10.1103/PhysRevApplied.15.064063](https://doi.org/10.1103/PhysRevApplied.15.064063). URL: <http://arxiv.org/abs/2103.07030> (visited on 02/22/2024).

[32] F. K. Wilhelm, S. Kirchhoff, S. Machnes, N. Wittler, and D. Sugny. *An introduction into optimal control for quantum technologies*. 2020.

[33] E. T. Holland, K. A. Wendt, K. Kravvaris, et al. "Optimal control for the quantum simulation of nuclear dynamics". *Phys. Rev. A* 101 (6 2020), p. 062307. DOI: [10.1103/PhysRevA.101.062307](https://doi.org/10.1103/PhysRevA.101.062307). URL: <https://link.aps.org/doi/10.1103/PhysRevA.101.062307>.

[34] A. Carignan-Dugas, D. Dahlen, I. Hincks, et al. "The Error Reconstruction and Compiled Calibration of Quantum Computing Cycles" (2023). arXiv:2303.17714 [quant-ph]. DOI: [10.48550/arXiv.2303.17714](https://doi.org/10.48550/arXiv.2303.17714). URL: <http://arxiv.org/abs/2303.17714> (visited on 05/02/2023).

[35] A. P. M. Place, L. V. H. Rodgers, P. Mundada, et al. "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds". en. *Nat Commun* 12.1 (2021). Number: 1 Publisher: Nature Publishing Group, p. 1779. DOI: [10.1038/s41467-021-22030-5](https://doi.org/10.1038/s41467-021-22030-5). URL: <https://www.nature.com/articles/s41467-021-22030-5> (visited on 01/24/2023).

[36] D. M. Tennant, L. A. Martinez, K. M. Beck, et al. "Low-Frequency Correlated Charge-Noise Measurements Across Multiple Energy Transitions in a Tantalum Transmon". *PRX Quantum* 3.3 (2022). Publisher: American Physical Society, p. 030307. DOI: [10.1103/PRXQuantum.3.030307](https://doi.org/10.1103/PRXQuantum.3.030307). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.3.030307> (visited on 01/24/2023).

[37] J. R. Johansson, P. D. Nation, and F. Nori. "QuTiP: An open-source Python framework for the dynamics of open quantum systems". en. *Computer Physics Communications* 183.8 (2012), pp. 1760–1772. DOI: [10.1016/j.cpc.2012.02.021](https://doi.org/10.1016/j.cpc.2012.02.021). URL: <https://www.sciencedirect.com/science/article/pii/S0010465512000835> (visited on 01/24/2023).

[38] J. R. Johansson, P. D. Nation, and F. Nori. "QuTiP 2: A Python framework for the dynamics of open quantum systems". en. *Computer Physics Communications*

References

184.4 (2013), pp. 1234–1240. DOI: [10.1016/j.cpc.2012.11.019](https://doi.org/10.1016/j.cpc.2012.11.019). URL: <https://www.sciencedirect.com/science/article/pii/S0010465512003955> (visited on 01/24/2023).

[39] X. Wu, S. Tomarken, N. A. Petersson, et al. “High-Fidelity Software-Defined Quantum Logic on a Superconducting Qudit”. en. *Phys. Rev. Lett.* 125.17 (2020), p. 170502. DOI: [10.1103/PhysRevLett.125.170502](https://doi.org/10.1103/PhysRevLett.125.170502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.170502> (visited on 01/05/2023).

[40] P. V. Klimov, J. Kelly, Z. Chen, et al. “Fluctuations of Energy-Relaxation Times in Superconducting Qubits”. *Phys. Rev. Lett.* 121.9 (2018). Publisher: American Physical Society, p. 090502. DOI: [10.1103/PhysRevLett.121.090502](https://doi.org/10.1103/PhysRevLett.121.090502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.090502> (visited on 04/24/2023).

[41] Z. Peng, D. Appelö, N. A. Petersson, et al. *Deterministic and Bayesian Characterization of Quantum Computing Devices*. arXiv:2306.13747 [quant-ph]. 2023. DOI: [10.48550/arXiv.2306.13747](https://doi.org/10.48550/arXiv.2306.13747). URL: [http://arxiv.org/abs/2306.13747](https://arxiv.org/abs/2306.13747) (visited on 07/08/2023).

[42] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. en. ISBN: 9780511976667 Publisher: Cambridge University Press. 2010. DOI: [10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667). URL: <https://www.cambridge.org/highereducation/books/quantum-computation-and-quantum-information/01E10196D0A682A6AEFFEA52D53BE9AE> (visited on 02/09/2023).

[43] Y. S. Weinstein, M. A. Pravia, E. M. Fortunato, S. Lloyd, and D. G. Cory. “Implementation of the Quantum Fourier Transform”. *Phys. Rev. Lett.* 86.9 (2001). Publisher: American Physical Society, pp. 1889–1891. DOI: [10.1103/PhysRevLett.86.1889](https://doi.org/10.1103/PhysRevLett.86.1889). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.86.1889> (visited on 02/16/2023).

[44] A. Peruzzo, J. McClean, P. Shadbolt, et al. “A variational eigenvalue solver on a photonic quantum processor”. en. *Nat Commun* 5.1 (2014). Number: 1 Publisher: Nature Publishing Group, p. 4213. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213). URL: <https://www.nature.com/articles/ncomms5213> (visited on 02/09/2023).

[45] A. Kandala, A. Mezzacapo, K. Temme, et al. “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets”. en. *Nature* 549.7671 (2017). Number: 7671 Publisher: Nature Publishing Group, pp. 242–246. DOI: [10.1038/nature23879](https://doi.org/10.1038/nature23879). URL: <https://www.nature.com/articles/nature23879> (visited on 02/16/2023).

[46] D. Wang, O. Higgott, and S. Brierley. “Accelerated Variational Quantum Eigensolver”. *Phys. Rev. Lett.* 122.14 (2019). Publisher: American Physical Society, p. 140504. DOI: [10.1103/PhysRevLett.122.140504](https://doi.org/10.1103/PhysRevLett.122.140504). URL:

References

<https://link.aps.org/doi/10.1103/PhysRevLett.122.140504> (visited on 02/16/2023).

[47] J. Tilly, H. Chen, S. Cao, et al. "The Variational Quantum Eigensolver: A review of methods and best practices". en. *Physics Reports* 986 (2022), pp. 1–128. DOI: [10.1016/j.physrep.2022.08.003](https://doi.org/10.1016/j.physrep.2022.08.003). URL: <https://www.sciencedirect.com/science/article/pii/S0370157322003118> (visited on 02/16/2023).

[48] Z. Liang, J. Cheng, H. Ren, et al. *PAN: Pulse Ansatz on NISQ Machines*. arXiv:2208.01215 [quant-ph]. 2022. URL: <http://arxiv.org/abs/2208.01215> (visited on 02/09/2023).

[49] V. Malkin, G. Shvets, and N. Fisch. "Fast compression of laser beams to highly overcritical powers". *Physical review letters* 82.22 (1999), p. 4448.

[50] R. Davidson. *Methods in nonlinear plasma theory*. Elsevier, 2012.

[51] Y. Shi, A. R. Castelli, X. Wu, I. Joseph, V. Geyko, F. R. Graziani, S. B. Libby, J. B. Parker, Y. J. Rosen, L. A. Martinez, et al. "Simulating non-native cubic interactions on noisy quantum machines". *Physical Review A* 103.6 (2021), p. 062608.

[52] S. Caldwell, N. Didier, C. Ryan, E. Sete, A. Hudson, P. Karalekas, R. Manenti, M. da Silva, R. Sinclair, E. Acala, et al. "Parametrically activated entangling gates using transmon qubits". *Physical Review Applied* 10.3 (2018), p. 034050.

[53] C. Huang, T. Wang, F. Wu, et al. "Quantum Instruction Set Design for Performance". *Physical Review Letters* 130.7 (2023). Publisher: American Physical Society, p. 070601. DOI: [10.1103/PhysRevLett.130.070601](https://doi.org/10.1103/PhysRevLett.130.070601). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.070601> (visited on 05/15/2024).

[54] B. Nachman, M. Urbanek, W. A. de Jong, and C. W. Bauer. "Unfolding quantum computer readout noise". *npj Quantum Information* 6.1 (2020), p. 84.

[55] N. Didier. "Flux control of superconducting qubits at dynamical sweet spots". *arXiv:1912.09416 [quant-ph]* (2019). arXiv: 1912.09416. URL: <http://arxiv.org/abs/1912.09416> (visited on 01/18/2022).

[56] E. C. Peterson, G. E. Crooks, and R. S. Smith. "Fixed-Depth Two-Qubit Circuits and the Monodromy Polytope". en. *Quantum* 4 (2020). arXiv:1904.10541 [quant-ph], p. 247. DOI: [10.22331/q-2020-03-26-247](https://doi.org/10.22331/q-2020-03-26-247). URL: <http://arxiv.org/abs/1904.10541> (visited on 06/22/2022).

[57] J. J. Wallman and J. Emerson. "Noise tailoring for scalable quantum computation via randomized compiling". *Phys. Rev. A* 94 (5 2016), p. 052325. DOI: [10.1103/PhysRevA.94.052325](https://doi.org/10.1103/PhysRevA.94.052325). URL: <https://link.aps.org/doi/10.1103/PhysRevA.94.052325>.

References

[58] A. Hashim, R. K. Naik, A. Morvan, et al. "Randomized compiling for scalable quantum computing on a noisy superconducting quantum processor". *arXiv:2010.00215* (2020).

[59] S. J. Beale, K. Boone, A. Carignan-Dugas, et al. *True-Q*. Version 2. 2020. DOI: [10.5281/zenodo.3945250](https://doi.org/10.5281/zenodo.3945250). URL: <https://doi.org/10.5281/zenodo.3945250>.

[60] J.-L. Ville, A. Morvan, A. Hashim, et al. "Leveraging randomized compiling for the quantum imaginary-time-evolution algorithm". *Phys. Rev. Res.* 4 (3 2022), p. 033140. DOI: [10.1103/PhysRevResearch.4.033140](https://doi.org/10.1103/PhysRevResearch.4.033140). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.4.033140>.

[61] A. Avtandilyan and W. Pogosov. "Optimal-order Trotter-Suzuki decomposition for quantum simulation on noisy quantum computers". *arXiv:2405.01131* (2024).

[62] Y. Shi, A. R. Castelli, X. Wu, I. Joseph, V. Geyko, F. R. Graziani, S. B. Libby, J. B. Parker, Y. J. Rosen, L. A. Martinez, and J. L. DuBois. "Simulating non-native cubic interactions on noisy quantum machines". *en. Phys. Rev. A* 103.6 (2021), p. 062608. DOI: [10.1103/PhysRevA.103.062608](https://doi.org/10.1103/PhysRevA.103.062608). URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.062608> (visited on 01/05/2023).

[63] V. E. Zakharov and L. A. Takhtadzhyan. "Equivalence of the nonlinear Schrödinger equation and the equation of a Heisenberg ferromagnet". *Theoretical and Mathematical Physics* 38.1 (1979), pp. 17–23.

[64] T. Hoffmann. "On the equivalence of the discrete nonlinear Schrödinger equation and the discrete isotropic Heisenberg magnet". *Physics Letters A* 265.1-2 (2000), pp. 62–67. DOI: [10.1016/S0375-9601\(99\)00860-9](https://doi.org/10.1016/S0375-9601(99)00860-9). URL: [https://doi.org/10.1016/S0375-9601\(99\)00860-9](https://doi.org/10.1016/S0375-9601(99)00860-9).

[65] V. E. Zakharov et al. "Collapse of Langmuir waves". *Sov. Phys. JETP* 35.5 (1972), pp. 908–914.

[66] J. M. Gambetta, F. Motzoi, S. T. Merkel, and F. K. Wilhelm. "Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator". *Phys. Rev. A* 83.1 (2011). Publisher: American Physical Society, p. 012308. DOI: [10.1103/PhysRevA.83.012308](https://doi.org/10.1103/PhysRevA.83.012308). URL: <https://link.aps.org/doi/10.1103/PhysRevA.83.012308> (visited on 02/17/2023).

[67] J. Chen, D. Ding, C. Huang, and Q. Ye. *Compiling Arbitrary Single-Qubit Gates Via the Phase-Shifts of Microwave Pulses*. *arXiv:2105.02398* [quant-ph]. 2022. URL: [http://arxiv.org/abs/2105.02398](https://arxiv.org/abs/2105.02398) (visited on 05/17/2023).

[68] F. Arute, K. Arya, R. Babbush, et al. *Nature* 574 (2019), p. 505.

[69] A. Arrasmith, A. Patterson, A. Boughton, and M. Paini. *Development and Demonstration of an Efficient Readout Error Mitigation Technique for use in NISQ Algorithms*. 2023. URL: <https://arxiv.org/abs/2303.17741>.

References

[70] J. J. Wallman and J. Emerson. "Noise tailoring for scalable quantum computation via randomized compiling". en. *Physical Review A* 94.5 (2016). ZSCC: 0000103 arXiv: 1512.01098, p. 052325. DOI: [10.1103/PhysRevA.94.052325](https://doi.org/10.1103/PhysRevA.94.052325). URL: <http://arxiv.org/abs/1512.01098> (visited on 03/11/2021).

[71] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio. "Error mitigation with Clifford quantum-circuit data". *Quantum* 5 (2021). arXiv:2005.10189 [quant-ph], p. 592. DOI: [10.22331/q-2021-11-26-592](https://doi.org/10.22331/q-2021-11-26-592). URL: <http://arxiv.org/abs/2005.10189> (visited on 04/17/2024).

[72] E. Rosenberg, T. Andersen, R. Samajdar, et al. "Dynamics of magnetization at infinite temperature in a Heisenberg spin chain". *Science* 384.6691 (2024). arXiv:2306.09333 [quant-ph], pp. 48–53. DOI: [10.1126/science.adl7877](https://doi.org/10.1126/science.adl7877). URL: <http://arxiv.org/abs/2306.09333> (visited on 04/24/2024).

[73] D. C. McKay, I. Hincks, E. J. Pritchett, et al. *Benchmarking Quantum Processor Performance at Scale*. 2023. URL: <https://arxiv.org/abs/2311.05933>.

[74] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta. "Efficient Z-Gates for Quantum Computing". *Physical Review A* 96.2 (2017). arXiv:1612.00858 [quant-ph], p. 022330. DOI: [10.1103/PhysRevA.96.022330](https://doi.org/10.1103/PhysRevA.96.022330). URL: <http://arxiv.org/abs/1612.00858> (visited on 12/12/2023).

[75] Y. Kim, A. Eddins, S. Anand, et al. "Evidence for the utility of quantum computing before fault tolerance". en. *Nature* 618.7965 (2023). Number: 7965 Publisher: Nature Publishing Group, pp. 500–505. DOI: [10.1038/s41586-023-06096-3](https://doi.org/10.1038/s41586-023-06096-3). URL: <https://www.nature.com/articles/s41586-023-06096-3> (visited on 06/14/2023).

[76] E. A. Sete, N. Didier, A. Q. Chen, et al. "Parametric-Resonance Entangling Gates with a Tunable Coupler". *Phys. Rev. Appl.* 16 (2 2021), p. 024050. DOI: [10.1103/PhysRevApplied.16.024050](https://doi.org/10.1103/PhysRevApplied.16.024050). URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.16.024050>.

[77] M. A. Rol, F. Battistel, F. K. Malinowski, et al. "Fast, High-Fidelity Conditional-Phase Gate Exploiting Leakage Interference in Weakly Anharmonic Superconducting Qubits". *Phys. Rev. Lett.* 123 (12 2019), p. 120502. DOI: [10.1103/PhysRevLett.123.120502](https://doi.org/10.1103/PhysRevLett.123.120502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.120502>.