

Characterizing and Mitigating Errors in Quantum Computers

by

Swarnadeep Majumder

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

Kenneth R. Brown, Supervisor

Jungsang Kim

Iman Marvian

Michael Gehm

Raphael C. Pooser

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Electrical and Computer Engineering
in the Graduate School of Duke University

2023

ABSTRACT

Characterizing and Mitigating Errors in Quantum
Computers

by

Swarnadeep Majumder

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

Kenneth R. Brown, Supervisor

Jungsang Kim

Iman Marvian

Michael Gehm

Raphael C. Pooser

An abstract of a dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in the Department of Electrical and Computer
Engineering in the Graduate School of Duke University

2023

Copyright © 2023 by Swarnadeep Majumder
All rights reserved except the rights granted by the
Creative Commons Attribution-Noncommercial Licence

Abstract

This thesis aims to present methods for characterizing and mitigating errors in quantum computers. We begin by providing a historical overview of computing devices and the evolution of quantum information. The basics of characterizing noise in quantum computers and the utilization of quantum control and error mitigation techniques to reduce the impact of noise on performance are also discussed. In the initial part of the thesis, we focus on a particularly detrimental type of time-dependent errors and derive theoretical limits of a closed-loop feedback based quantum control protocol for their mitigation. Two different protocols, one suitable for fault-tolerant systems and another for near-term devices, are presented and their performance is demonstrated through numerical simulations. Additionally, we explore the mitigation of coherent noise at the circuit level through the use of the hidden inverses protocol with results from experiments conducted at Duke University, Sandia National Laboratories, and IBM. Finally, we propose a scalable error characterization procedure for large quantum systems, which is tested through numerical simulations to highlight its sensitivity to various sources of noise. Crucially, this protocol does not require access to ideal classical simulation of quantum circuits unlike other benchmarks such as quantum volume or cross entropy benchmarks.

Contents

Abstract	iv
List of Figures	viii
Acknowledgements	x
1 Introduction	1
1.1 Characterizing noise in quantum computers	4
1.2 Mitigating noise in quantum computers	6
1.2.1 Quantum Control	6
1.2.2 Quantum Error Mitigation	7
2 Real time calibration with spectator qubits	9
2.1 Introduction	9
2.2 Results	12
2.2.1 Theoretical limits	12
2.2.2 Application to magnetic field noise	19
2.2.3 Application to laser beam instability	22
2.3 Discussion	28
2.4 Methods	30
2.4.1 Error model	30
2.4.2 Magnetic field noise	33
2.4.3 Laser beam instability	34

2.5	Key Observations	35
2.6	Single shot correction theory	37
2.7	Low overhead spectator qubit algorithm	40
2.8	Results	41
2.9	Discussion	42
3	Circuit level error mitigation with hidden inverses	46
3.1	Introduction	46
3.2	Hidden inverses protocol to mitigate coherent errors	48
3.3	Experimental implementation of quantum circuits on a trapped ion processor	49
3.4	Demonstration of hidden inverses as an error characterization experiment with single qubit gates	52
3.4.1	Experimental design	52
3.4.2	Classical simulation and noise model	54
3.4.3	Result	54
3.5	Hidden inverses as an error mitigation protocol for applications: variational quantum eigensolvers	55
3.5.1	Hamiltonian and Ansatz	55
3.5.2	Randomized Compiling	57
3.5.3	Purification	57
3.5.4	Noise Injection	58
3.6	Key observations	60
3.7	Hidden inverses in superconducting hardware	63
3.7.1	Gate inversion on superconducting based architectures	64
3.7.2	Application to Variational Quantum algorithms	66
3.7.3	Discussion	68
3.8	Conclusions	73

4	Characterizing large scale quantum systems	75
4.1	Estimating average fidelity of cycles	77
4.2	Algorithmic modification in the presence of coherent error	78
4.3	Numerical simulations and Experimental implementations	78
4.3.1	Noise model for numerical simulations	78
4.3.2	Testing for spatial in-homogeneity	84
4.4	Derivation of the fit model	84
4.4.1	Estimating purity from Swap test	84
4.4.2	Error scrambling under random circuits	84
4.4.3	Estimating depolarizing probability and fidelity	85
5	Conclusion	86
	Biography	104

List of Figures

2.1	General feedback loop setup for spectator qubits.	12
2.2	Diagrams of two spectator qubit applications.	20
2.3	Classical field error with and without spectator qubits.	21
2.4	Beam instability errors, with and without spectator qubits.	24
2.5	Example of a frustrated spectator qubit scheme.	26
2.6	Control landscapes for beam instability errors.	27
2.7	Single shot spectator qubit protocol	40
2.8	Diagrammatic view of Algorithm 1	42
2.9	Stationary Regime	44
2.10	Non-stationary Regime	44
2.11	Effective error parameter	45
3.3	Illustration of noise characterization using single qubit hidden inverses.	52
3.5	Hidden inverses as an error characterization tool	55
3.6	Error parameter drift	56
3.7	VQE tapered ansatz	57
3.8	Comparison of HI and RC as mitigation protocols for coherent error	61
3.9	Comparison of HI and RC as mitigation protocols for cooling loop noise	62
3.10	Pulse profile for CX^{-1} and verification:	65

3.11	Effective noise scaling on CX gate - IBMQ-Montreal:	66
3.12	The hidden inverse in the UCC3 ansatz	67
3.13	Energy ground state estimation for H_2 using VQE	68
3.14	Comparing the effect of simulated noise in VQE optimization landscape and the result of applying randomized compiling: .	72
3.15	Effect of simulated time-dependent coherent noise in VQE optimization landscape and result of applying HI	73
4.1	Estimating error per circuit layer.	76
4.2	Purity plotted against the number of layers for a noise model with single qubit and two qubit depolarizing noise.	79
4.3	Purity plotted against the number of layers for a noise model with purely coherent error.	80
4.4	Purity plotted against the number of layers for a noise model with purely coherent error with twirling	81
4.5	Comparison of purity plotted against the number of layers for a noise model with a cross-talk and one without.	82
4.6	Effect of noisy swap test on the performance of the protocol.	83
4.7	Swap test circuit	85

Acknowledgements

Graduate studies present a unique and challenging journey, and I am grateful for the support and guidance of the people who have made my experience at Duke University unforgettable.

First and foremost, I would like to express my gratitude to my advisor, Ken Brown, who provided me with invaluable guidance, encouragement, and wisdom throughout my studies. Ken struck the perfect balance between giving me the freedom to pursue my research goals and keeping me grounded with critical questions. His expertise and insights in the field of quantum information science have shaped me into a better researcher and person. I am grateful for his unwavering support and mentorship, and I hope to one day have the same impact on others as he has had on me.

I would also like to extend my sincere thanks to my committee members, Jungsang Kim, Iman Marvian, Michael Gehm, and Raphael Pooser, for their collaborations, valuable feedback and suggestions throughout my research. Their expertise and dedication have contributed greatly to the quality of my work.

I am thankful for the encouragement and support from my former professors, including Shanshan Rodriguez, Leo Rodriguez, Reinhold Ludwig, P.K. Aravind, William Baller, and Michael Johnson from Worcester Polytechnic Institute. These individuals have had a profound impact on my academic and personal growth.

I would like to express my appreciation to the graduate students, postdoctoral as-

sociates, and research scientists in our lab who made my time at Duke unforgettable. I am grateful to Leo Castro, who was a post-doc in my first year, for his patience and help in guiding me as a scientist. I also have fond memories of working with the rest of the control team Mingyu Kang, Travis Hurant, and James Leung. I am thankful for the rest of our group, including Natalie Brown, Dripto Debory, Muyuan Li, Shilin Huang, Evan Reed, Brad Bondurant, Jacob Whitlow, Swapnil Patel, Aniket Dalvi, Chao Fang, Leon Riesebo, Ely Novakoski, Bichen Zhang, Gloria Jia, Ke Sun, Omid Khosravani, Kavyashree Ranawat, Phoebe Chen, Balint Pato, Sagik Saha and Stephen Crain. I have always enjoyed many conversations with Mike Newman, Ye Wang, Jyothi Saraladevi, Abhinav Anand, and Jonathan Baker.

Collaboration is key to success in any scientific field, and I am grateful for the opportunity to work with talented undergraduate students at Duke, Faris Sabhi, Catherine Liang, and Raymond Lin, as well as my external collaborators Ryan Bennink, Vicente Leyton and Titus Morris from Oak Ridge National Laboratory, Christopher Yale from Sandia National Laboratory, and Pranav Gokhale from Inflation. I also had a productive summer internship at IBM Quantum, where I had the pleasure of working with Zlatko Minev, Derek Wang, Haimeng Zhang, Alireza Seif, Mirko Amico, Ben Roasnd, Kunal Sharma and Sona Najafi.

I am grateful for the friendships I have made. I would like to thank Shankha Dandapat and Aritrajit Gupta, whom I met in high school and have remained close friends with until today. I would also like to thank Alex Rawley, Kushagra Kumar, Siddhant Pandey, Shanshan Xie, Kayleah Griffen, Andrew Aberdale, and Matt Carnein for making my college life enjoyable. I am thankful for my friends in Durham, including Tamoghna Saha, Natalie Rozman, Ava Hurlock, Jessica and Aaron Centers, Nico Lozada, Kavin Sivakumar, Janet Chen, Brittani and Dillon Huegen, and Reed Hodges.

I would like to thank the National Science Foundation for their QISE-NET fellow-

ship. QISE-NET has allowed me to work closely with Dr. Pooser and the rest of the Oak Ridge team. Further, I would also like to show my gratitude towards Department of Defense Multidisciplinary University Research Initiative (MURI), National Science Foundation (NSF), Department of Energy (DOE) and Intelligence Advanced Research Projects Activity (IARPA) for funding my studies and research projects throughout the years.

Finally, I would like to thank my family, especially my parents Sankar Majumder and Sudipta Majumder, for their love, support, and sacrifice. Their unwavering belief in me has been the foundation of my success, and I am forever grateful for their encouragement and guidance. I would also to thank my brother Suvodeep Majumder who has been always there for me.

Introduction

Computing devices [Bec+03] have come a long way since their inception and they have played a transformative role in many areas of life. Early computing devices such as the abacus and the slide rule were used for basic mathematical calculations. For example, the invention of the mechanical calculator in the 19th century made it possible to perform more complex calculations, and the first electronic computer, the Electronic Numerical Integrator and Computer (ENIAC) [ME47], was developed during World War II and was used for military calculations such as encryption breaking.

In the 1950s and 1960s, mainframe computers were introduced, which were large and expensive machines used primarily by businesses and government institutions. However, the invention of the microprocessor in the 1970s paved the way for the development of personal computers, which were smaller and more affordable.

The 1980s saw the rise of the personal computer with the introduction of the IBM PC and Apple's Macintosh. The 1990s brought the development of the Internet, which revolutionized the way we communicate and access information. The 2000s saw the rise of mobile computing devices such as smartphones and tablets, which have become an integral part of our daily lives. Today, computing devices continue to evolve and include virtual and augmented reality devices as well as the Internet of Things (IoT) devices.

Moreover, computing has had a significant impact on scientific discoveries in a

variety of fields. In physics, computing has been used to simulate and analyze the behavior of subatomic particles [AT17], thereby helping to further our understanding of the fundamental nature of matter and energy. In chemistry, computing has been used to predict the properties and behavior of molecules, which allowed for the design of new drugs and materials [AW59; Rah64].

In biology, computing has been used to analyze and interpret large amounts of genetic data, which has led to new insights into the genetic basis of disease and the development of personalized medicine [KD18]. In medicine, computing has been used to develop new imaging techniques, such as magnetic resonance imaging (MRI) and computed tomography (CT) scans, which have revolutionized the diagnosis and treatment of diseases [MG75].

In earth science, computing has been used to model and predict weather patterns, sea levels, and climate change [Ben+82]. In astronomy, computing has been used to process and analyze data from telescopes and satellites, leading to new discoveries about the origins of the universe and the nature of other celestial bodies [Col+19].

Overall, computing has played a crucial role in scientific discoveries by allowing scientists to analyze and interpret large amounts of data, model complex systems, and make predictions. It has also greatly increased the speed and accuracy of scientific research and has led to new discoveries that would have been impossible without the use of computing technology.

Quantum mechanics is a branch of physics that deals with the behavior of matter and energy at the atomic and subatomic level. It was first developed in the early 20th century as a way to explain phenomena that could not be explained by classical mechanics, such as the behavior of electrons in atoms [Bor26].

The development of quantum mechanics began in 1900 with Max Planck's explanation of the black-body radiation problem. He proposed that energy is not continuous but instead is quantized meaning that it exists in discrete units. This

idea was later developed by Albert Einstein, who used it to explain the photoelectric effect [Fey63].

In the 1920s, a series of experiments were performed that further supported the idea of quantized energy [GS22]. In 1924, Louis de Broglie proposed that particles, such as electrons, can also exhibit wave-like behavior. This idea was confirmed by experiments performed by Davisson and Germer in 1927 [DG28].

The development of quantum mechanics was also heavily influenced by the work of Niels Bohr [Kra12], who proposed that electrons in atoms can only occupy certain discrete energy levels, and by Erwin Schrödinger [Sch26] and Werner Heisenberg, who formulated the wave equation and the matrix mechanics respectively to describe the behavior of nature at the level of atoms.

In the 1930s and 1940s, the field of quantum mechanics was further developed by the work of Paul Dirac [Dir30; Dir28], Enrico Fermi, and others, who developed the theory of quantum electrodynamics (QED) to describe the behavior of electrons and photons. John Clauser, Alain Aspect and Anton Zeilinger performed experiments with entangled states to test of Bell inequalities [ADR82; Cla+69]. This evidence that nature at small scale follows a different kind of physical laws than what we experience in our daily lives started the quantum revolution. The first quantum revolution gave us transistors, atomic clocks, lasers which form the bedrock of modern computing, navigation and communication [BB48; EP55]. Although these technologies brought about radical advancements in scientific discoveries, scientists (namely Richard Feynman [Fey82] and Yuri Manin[Man80]) in 1980s realized a computer based on quantum information would be more efficient in solving problems that are fundamentally quantum mechanical in nature than computers based classical logic. While the idea that quantum computers can efficiently solve quantum simulation problems seems intuitive, researchers also discovered their advantage in different computational tasks unrelated to simulating quantum systems. These include

Charles Bennett’s cryptography protocols [BB84] to enhance storage and transmission of sensitive information, oracle problems including Deutsch’s algorithm [Deu92], the Bernstein–Vazirani algorithm [BV97] and finally Peter Shor’s famous polynomial time algorithm for factoring large numbers into its primes [Sho94]. In the 1990s, discoveries in algorithms and theories were accompanied by remarkable advancements in hardware to demonstrate highly controllable quantum systems [Mon+95; CZ95]. These developments together formed the bedrock of the second quantum revolution, quantum information on quantum hardware.

In the field of quantum computing, handling noise presents a major challenge to the advancement of this technology. This is because noise significantly reduces the capabilities of a quantum computer. This dissertation aims to investigate methods of characterizing noise and mitigating its impact on performance. In the first chapter, we provide an overview of the fundamental principles of quantum information theory and introduce key concepts such as quantum characterization, verification, validation, and quantum error mitigation. In the second chapter, we propose real-time feedback-based control protocols using spectator qubits to counteract time-dependent noise. In the third chapter, we investigate the application of hidden inverses, a circuit-level error mitigation technique to correct for coherent errors and demonstrate its efficacy through experiments performed at Sandia National Lab and IBM. In the fourth chapter, we propose a scalable error characterization procedure to benchmark large quantum computers without relying on classical simulation. Finally, we summarize our findings and conclusions in the fifth chapter.

1.1 Characterizing noise in quantum computers

The central components of building a functional quantum computer are qubits, which are physical systems capable of storing quantum information primitives, and the means to manipulate this information through quantum gates. In order to extract

classical information from the final quantum states, a readout or measurement is also necessary. However, deviations from the ideal mathematical counterparts of these physical operations can result in “noise”, which can be further divided into state preparation and measurement error (known collectively as SPAM error) and gate error.

To address these errors, the field of quantum characterization, verification, and validation (QCVV) has emerged. Tomography is considered the most comprehensive method for evaluating the quality of a quantum state or operation, as it provides a framework for reconstructing an unknown quantum state or process based on experimental data. However, for more than three qubits, this method becomes impractical due to the high number of samples required. Alternative approaches, such as randomized benchmarking [MGE12; EAŽ05], provide a way to extract partial information about quantum processes, such as the average gate fidelity or process fidelity, to quantify the difference between the ideal quantum operation and the noisy implementation. Other methods, such as quantum volume [Cro+19] and cross entropy benchmarks [Boi+18], have been proposed to characterize the performance of large-scale systems. However, they require access to ideal classical simulations, which limits the size of the systems that can be evaluated.

In Chapter 2, the attention is directed towards time-dependent errors that can cause significant damage, and a theoretical limit is established for closed-loop quantum control techniques to counteract these errors using spectator qubits. A protocol to reach this limit for realistic noise models is described and demonstrated through numerical simulation. Additionally, in collaboration with Faris Sabhi and Dr. Iman Marvian, we developed a spectator qubit protocol that is more suited for near-term devices. Chapter 4 introduces a protocol for characterizing large-scale quantum systems, which does not require classical simulations, unlike the quantum volume or cross entropy benchmarks.

1.2 Mitigating noise in quantum computers

Due to the inherent noise in physical systems, quantum computers can be unreliable without quantum error correction. To overcome this issue and extract meaningful results from computations, effective methods for dealing with noise are necessary. There are numerous proposals for limiting the impact of noise, which can be broadly classified into two categories: (1) quantum control techniques at the hardware level and (2) quantum error mitigation at the circuit or algorithmic level.

1.2.1 *Quantum Control*

Quantum control involves the active management of the evolution of quantum systems. It has been applied in a variety of fields, including atomic physics, optics, molecular chemistry, and quantum biology. At the heart of quantum computing is the control of a quantum system through the manipulation of the properties of an external electromagnetic field. By adjusting these properties, quantum gates specific to a particular system can be designed. By gaining a thorough understanding of the underlying noise, the parameters can be further optimized to make the quantum operations resistant to noise. The process of learning the system's properties and adapting the control protocol is called calibration.

Several techniques are commonly used in gate design, including composite pulses [Lev86; Wim94; BHC04; TV11; Vit11; Ban+13; Gen+14; LYC14; Kab+14; MB14; LYC16; CRR16] such as SK1 or BB1, which are widely used to correct systematic errors in gates. Entangling operations are made to be robust against various noises using frequency modulation [Leu+18], amplitude modulation [Roo08], phase modulation [Mil+20], and batch optimization [Kan+21]. Additionally, gate idling errors can be minimized through the use of dynamical decoupling pulses [VKL99; VK03; RHC10; SÁS12; QL13; Pok+18]. By combining these protocols, single qubit gate

fidelities of over 99.99% and two qubit gate fidelities of over 99.9% [Bal+16] have been achieved.

However, in the presence of time-dependent errors, the performance of these protocols may deteriorate over time as the underlying noise changes. To address this issue, in chapter 2, we describe a closed-loop feedback control protocol that allows for the re-calibration of a quantum computer.

1.2.2 *Quantum Error Mitigation*

The ultimate goal in the field of quantum computing is to develop an error-corrected and fault-tolerant quantum computer [AB97; Kit97; KLZ98; Got98; Kni05; AC07; Ter15]. However, implementing quantum error correction requires a substantial qubit overhead and operations that surpass a certain threshold. Despite recent advancements in experimental error correction, it is unlikely that these techniques will be applicable in practical settings in the near future. As a solution, quantum error mitigation protocols [MYB19; Bon+18; MZO20; Che+19; KB20; Str+20; Cza+20; ZG20] have been developed to reduce the impact of noise at the algorithmic level. The most commonly used protocols aim to reduce the bias in the expectation value of observables with a low to moderate overhead. This is achieved by running and combining the results from multiple circuits. Two examples of such error mitigation protocols are Probabilistic Error Cancellation (PEC) [Ber+22; TBG17] and Zero Noise Extrapolation (ZNE) [TBG17]. While these protocols minimize the impact on the ensemble of circuits, the noise level in individual circuits remains unchanged. Chapter 3 introduces a circuit-level error mitigation technique called "Hidden Inverses" that enhances the quality of individual circuits. The trapped ion experiments described in this chapter were conducted at Sandia National Laboratory by Dr. Christopher Yale, while the superconducting experiments were performed at IBM quantum hardware by Dr. Vicente Leyton-Ortega of Oak Ridge National Lab.

Our study shows that the use of hidden inverses can lead to improved performance of variational algorithms in both trapped ion and superconducting quantum computers.

Real time calibration with spectator qubits

The majority of this chapter is taken from Ref [MCB20]:

Swarnadeep Majumder, Leonardo Andreta de Castro, and Kenneth R. Brown. “Real-time calibration with spectator qubits”. *npj Quantum Information* 6 (1 Feb. 2020). doi: 10.1038/s41534-020-0251-y.

2.1 Introduction

One of the key challenges in constructing a quantum computer is keeping the error rate under an acceptable threshold, which will be a requirement even for future fault-tolerant quantum computation [AB97; Kit97; KLZ98; Got98; Kni05; AC07; Ter15]. The optimal control strategy for each quantum gate depends on the parameters that characterize the underlying error channel \mathcal{E} . There has been an increasing interest in tailoring control strategies to the error channel, such as variability-aware qubit allocation and movement [TQ19], optimal quantum control using randomized benchmarking [Kel+14], robust phase estimation [KLY15], noise-adaptive compilation [Mur+19], and quantum error-correcting codes designed for biased noise [NP13; TBF18; Tuc+19; Li+19].

Although an initial calibration may be sufficient for simpler devices, a fully functional quantum computer will have to deal with the possibility of assessing changes in the error parameters in real time. Many reduction techniques have been proposed for errors that vary slowly in time, such as composite pulses [Lev86; Wim94; BHC04;

TV11; Vit11; Ban+13; Gen+14; LYC14; Kab+14; MB14; LYC16; CRR16], optimal control [MCF07; Gra+07; Sch+11; Mac+18], dynamical decoupling [VKL99; VK03; RHC10; SÁS12; QL13; Pok+18], and dynamically corrected gates [KV09b; KV09a; GUB12]. In this work, we analyze the use of a subset of qubits – called spectator qubits – to perform real-time recalibration.

Spectator qubits probe directly the sources of error and thus do not need to interact with the data qubits, so they can be distinguished from ancilla qubits used for syndrome extraction [Sho95; Ter15] in quantum error correction. As long as the error channel of the spectator qubits is correlated to the error channel of the data qubits, it is possible to estimate \mathcal{E} by measuring the spectators. Although sensor networks [Qia+19], machine learning techniques [Mav+17; GB18], and even spectator qubits [Gup+19] have been proposed to keep track of error parameters that vary in space or time, more often than not these techniques are not suitable for real-time calibration because of how long it takes to extract useful information about the error parameters from the experimental data. Here we describe the complete feedback loop between the information extracted from the spectator qubits and the recalibration of the control strategy on the data qubits, estimating how this information can positively impact the control protocol. When the necessity for feedback is taken into account, acquisition of information via the spectator qubits has to be sufficiently fast such that the rate of errors in the data qubit does not exceed the rate at which the parameters are being estimated. Such feedback schemes could in principle deal with general classes of errors, but in this work we will limit our discussion to particularly damaging coherent errors.

We illustrate the difficulty of using feedback against coherent errors with a simple example. Consider constant overrotations around the x -axis characterized by the error parameter θ . If the error rate is the same as the rate of acquisition of information, the estimate of θ after N overrotations will have an imprecision proportional

to $N^{-1/2}$. For this reason, any attempt to correct the error with the inverse unitary will result in an extant error that still grows with $O(N^{1/2})$:

$$e^{iN\theta\sigma_x} e^{-iN[\theta+O(N^{-1/2})]\sigma_x} = e^{iO(N^{1/2})\sigma_x}. \quad (2.1)$$

This kind of difficulty is common to coherent errors in general, but can be contained with the help of quantum control techniques that reduce the speed with which the errors accumulate in the data qubits. Other ways of balancing the increase of errors with a sufficient fast acquisition of information are: usage of different species of qubits for data and spectators, and application of different rates of measurements in order to reach the Heisenberg limit. Here, we will focus on the first strategy of making the data qubits less sensitive to errors via control strategies, and leave the other methods to be explored in future work.

In this work, we propose that real-time calibration with spectator qubits can in principle improve the fidelity of any system undergoing coherent errors, as long as: (1) the information available to the spectator qubits is sufficient to keep track of the rate of change of the error parameters; and (2) we have a quantum control method capable of sufficiently suppressing the speed with which the coherent errors accumulate in the data. The general setup that we will consider is the one illustrated in Fig. 1, where spectator qubits embedded in the same architecture as the data qubits are measured periodically to determine the error profile. The errors might change in space and time, but as long as they have some correlation between neighboring qubits, we can extrapolate an error field that contains information about the error profile in the data qubits. A classical apparatus then uses the information about this profile and about the results of the spectator measurements to decide the best control strategy. Here, we present the theoretical framework for studying multipartite systems composed of spectator qubits and data qubits in presence of coherent errors, propose some applications, and present analytical and numerical simulations of the performance of

the spectator qubits.

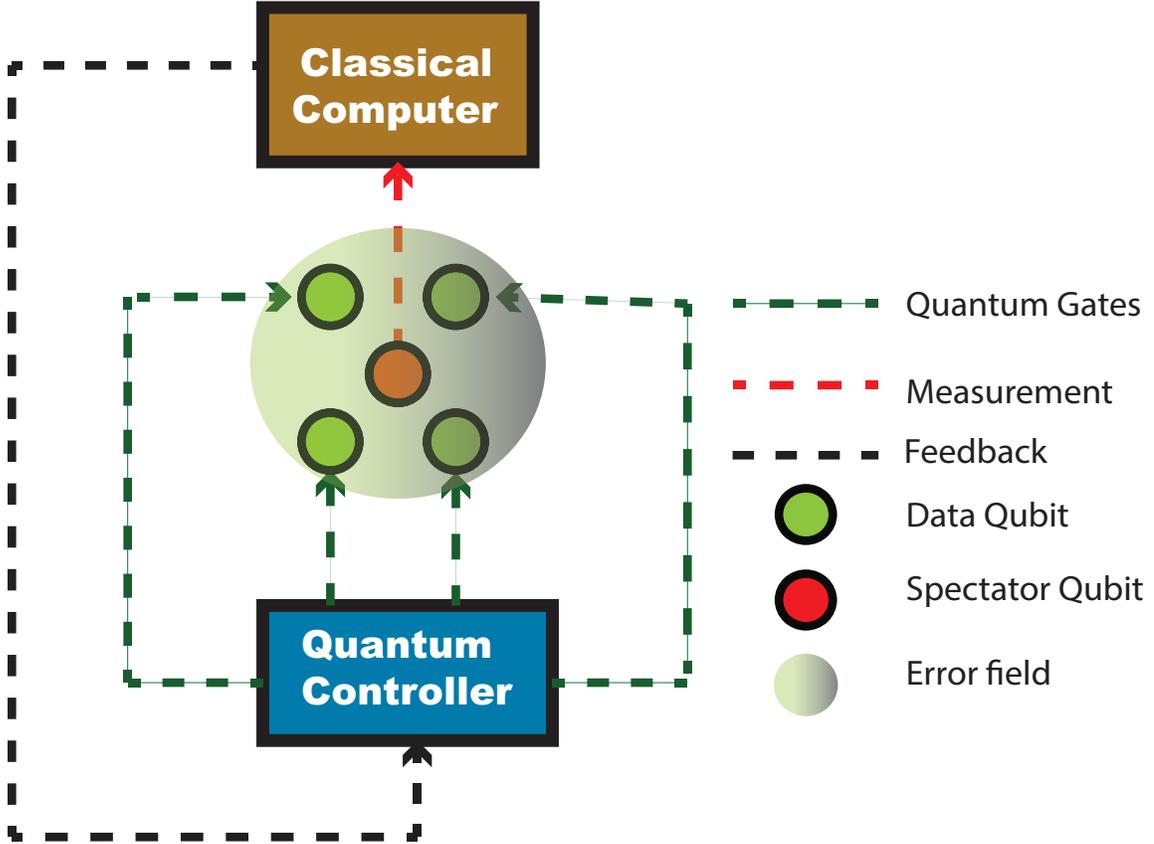


Figure 2.1: **General feedback loop setup for spectator qubits.** Information from the error field is acquired via measurements of the spectator qubits (red). This is analysed by a classical apparatus, which then updates the optimal control strategy for the data qubits (green).

2.2 Results

2.2.1 Theoretical limits

The purpose of spectator qubits is to obtain information about an error parameter while its value slowly drifts. This problem can be approximated by a setting where the error parameter is assumed to be fixed within a series of time slots. After the end of each time slot, the parameter changes according to a probability distribution.

In this work, we will assume that the drift is sufficiently slow that this error

parameter does not change significantly during a measurement. In this limit, the duration of the measurement process does not affect the result of the measurement, and we can assume it to be instantaneous. However, the number N of measurements that can be performed in the period of time where the error parameters remain fixed is still limited.

The precision with which we can learn about the drift from this series of N measurements is limited by the Cramér–Rao bound [BC94]. Calling our imprecision $\delta\vartheta$, the Cramér–Rao bound has the form:

$$\langle |\delta\vartheta| \rangle \geq \frac{1}{\sqrt{N f_\theta}}, \quad (2.2)$$

where f_θ is the Fisher information about the error parameter θ available to the system between each measurement. The error parameter θ could in principle represent any kind of information about the error, but in the case of coherent errors it will often mean an overrotation angle. Still in the context of coherent errors, if the system can be represented by a pure-state density matrix ρ , the Fisher information takes the specific form [BC94]:

$$f_\theta = 4\text{Tr} \left\{ \rho \left(\frac{\partial \rho}{\partial \theta} \right)^2 \right\} = (1 - \langle \hat{\mathbf{n}} \cdot \boldsymbol{\sigma} \rangle^2), \quad (2.3)$$

where we are expressing the unitary error operator in terms of the rotation axis $\hat{\mathbf{n}}$ as $U = e^{-i(\hat{\mathbf{n}} \cdot \boldsymbol{\sigma})\theta/2}$, and the expectation value is taken for the initial state of the system.

While the imprecision decreases with $1/\sqrt{N}$, there are other resources that are able to improve our estimate faster than that. Increasing the number of overrotations L between each measurement reduces the imprecision by $1/L$ – the so-called Heisenberg scaling – a kind of precision that can also be achieved by using entangled qubits [Win+92; Bol+96]. These schemes increase the achievable precision by

increasing the Fisher information between each measurement. However, the scaling by the number of measurements will remain the same, $1/\sqrt{N}$.

Calling the error parameter during the k th time slot θ_k , the optimal correction scheme against the coherent error $U(\theta_k)$ would be to simply apply its Hermitian conjugate, $U^\dagger(\theta_k)$. Our actual strategy might not be as good as the ideal one, so this description will give us an upper limit for the possible recovery. Moreover, if we only have access to an estimate $\hat{\theta}_k$ of θ_k , our strategy will be limited by the amount of precision we can achieve in obtaining this estimate. Although it is possible that our measurement scheme will not be able to saturate the Cramér–Rao bound, this will nevertheless provide an upper limit to how much precision we can achieve. An example of a situation where the Cramér–Rao bound is not saturated occurs when we attempt to apply Robust Phase Estimation (RPE), a procedure that prescribes doubling the number of gates before each measurement [Rud+17], resulting in a precision that achieves the Heisenberg scaling [KLY15] if there are no time constraints. With time constraints, this scaling is not always true. Given this risk of underperformance, and the fact that robustness of RPE to time-dependent errors is not well-known [Mei+19], in this work we will only use a limited number of gate repetitions before measurement and concentrate our analysis on the scaling $1/\sqrt{N}$ that comes from varying the number of measurements N .

Regardless of how we obtain the estimate $\hat{\theta}_k$, once we have a reliable value for it, the best possible evolution after the optimal correction strategy is represented by the effective unitary $V(\phi)$:

$$V(\phi) = U^\dagger(\hat{\theta}_k)U(\theta_k), \quad (2.4)$$

where ϕ is an effective error parameter that depends essentially on the difference between θ_k and $\hat{\theta}_k$. Using this, we can estimate the best possible process fidelity for

a given ϕ :

$$F(\phi) = \left| \frac{\text{Tr} \{V^\dagger(\phi=0)V(\phi)\}}{\text{Tr} \{\mathbb{I}\}} \right|^2, \quad (2.5)$$

which is proportional to the average fidelity [PMM07]. Here, $\phi = 0$ corresponds to perfect knowledge of the error parameter, allowing a perfect evolution of the system. Additionally invoking the fact that the fidelity for a coherent error should be a continuous function of the angle, we can expand this fidelity as a power series around $F(\phi = 0)$:

$$F(\phi) = \sum_{n=0}^{\infty} \phi^n F^{(n)}(0), \quad (2.6)$$

where we are using the following notation for the n th derivative of the fidelity:

$$F^{(n)}(x) = \left. \frac{d^n F}{d\phi^n} \right|_{\phi=x}. \quad (2.7)$$

By our choice of ϕ , the point $\phi = 0$ corresponds to perfect knowledge of the error, making the fidelity take its maximum value. Therefore, $F(\phi = 0) = 1$, $F^{(1)}(\phi = 0) = 0$, and $F^{(2)}(\phi = 0) < 0$. Expanding the expression up to the second order and representing the extra terms as a Lagrange remainder, we find:

$$F(\phi) = 1 + \frac{1}{2}\phi^2 F^{(2)}(0) + \frac{1}{2} \int_0^\phi dx (x - \phi)^2 F^{(3)}(x). \quad (2.8)$$

Suppose ϕ_n is the effective error parameter if we do not update our estimates using the spectator qubits, whereas ϕ_s is the effective error parameter if we use spectator qubits to acquire information. The feedback loop will be successful if the fidelity obtained using spectator qubits is, on average, superior to the fidelity that we obtain without using them. Therefore, we want to satisfy:

$$\langle F(\phi_s) \rangle > \langle F(\phi_n) \rangle, \quad (2.9)$$

which in Eq. (2.8) is equivalent to:

$$\langle \phi_s^2 \rangle + \frac{1}{F^{(2)}(0)} \left\langle \int_0^{\phi_s} dx (x - \phi_s)^2 F^{(3)}(x) \right\rangle < \langle \phi_n^2 \rangle + \frac{1}{F^{(2)}(0)} \left\langle \int_0^{\phi_n} dx (x - \phi_n)^2 F^{(3)}(x) \right\rangle. \quad (2.10)$$

While it is possible to satisfy the necessary condition above even in situations where our spectators have not helped to decrease the effective error parameters, usually we will want to further impose that:

$$|\phi_s| < |\phi_n|, \quad (2.11)$$

a condition that is ultimately limited by the Cramér–Rao bound. However, while we try to meet this condition, we also should make sure that the errors do not increase excessively. This can be translated into a second set of conditions that are not necessary, but are sufficient to satisfy inequality (2.10) when we impose (2.11). These conditions simply say that the higher-order terms of the expansion must be negligible in comparison to the second-order terms that originate condition (2.11):

$$\langle \phi_s^2 \rangle \gg \left| \frac{1}{F^{(2)}(0)} \left\langle \int_0^{\phi_s} dx (x - \phi_s)^2 F^{(3)}(x) \right\rangle \right|, \quad (2.12)$$

$$\langle \phi_n^2 \rangle \gg \left| \frac{1}{F^{(2)}(0)} \left\langle \int_0^{\phi_n} dx (x - \phi_n)^2 F^{(3)}(x) \right\rangle \right|. \quad (2.13)$$

Do notice that, as sufficient but not necessary conditions, we do not need to respect (2.12) and (2.13) in order to obtain satisfactory results with spectator qubits. However, if the feedback loop with the spectators is not improving the fidelity of the system despite condition (2.11) being met, then changing the scheme so that (2.12) and (2.13) are satisfied will suffice to make the strategy work.

It is worth remarking that this kind of analysis that requires a Taylor expansion around $\phi = 0$ is not adequate for metrics that do not have a derivative at $\phi = 0$,

such as the diamond norm, which is defined as:

$$\diamond_{\phi} \equiv \frac{1}{2} \max_{\rho} (\| (V(\phi = 0) \otimes 1)[\rho] - (V(\phi) \otimes 1)[\rho] \|_1). \quad (2.14)$$

For this reason, we opted to use the process fidelity as defined in Eq. (2.5) in our analyses. In any case, as there is a one-to-one correspondence between the diamond norm and the process fidelity for coherent errors, an improvement in one metric translates into an improvement in the other one as well.

Finally, let us consider a simple example of a situation where forcing conditions (2.12) and (2.13) to be respected also makes the complete feedback loop to function. Suppose a data qubit and a spectator qubit simultaneously suffer the same kind of overrotation, $e^{i\phi X}$. In this situation, the necessary condition (2.9) takes the following exact form after N measurements and overrotations:

$$\cos^2(N\phi_s) > \cos^2(N\phi_n). \quad (2.15)$$

If a measurement is performed after each overrotation, the best average estimate we can find for ϕ is given by $1/(2\sqrt{N})$, according to the Cramér–Rao bound and Eq. (2.3). A simple way of satisfying (2.15) is by restricting the arguments of the cosines to the interval $[-\pi/2, \pi/2]$ and then imposing condition (2.10). However, ϕ_s will only be smaller than ϕ_n if we perform a number of measurements that is sufficiently large to have a clear estimate of ϕ_n – i.e., if $1/(2\sqrt{N}) < |\phi_n|$. By this point, however, the increase in N may have brought the angles out of the $[-\pi/2, \pi/2]$ range, in which case the necessary condition (2.15) may no longer be satisfied by simply reducing the value of ϕ_s . In particular, this translates into a violation of the sufficient condition (2.12). Instead, we have (see Section I of the Supplementary Material for details of the derivation):

$$\left| \frac{1}{F^{(2)}(0)} R_2(\phi_s) \right| = \left| \frac{1}{4N} + \frac{\cos(\sqrt{N}) - 1}{2N^2} \right|, \quad (2.16)$$

whose right-hand side approaches $\phi_s \sim 1/\sqrt{2N}$ as N grows, which causes a violation of the condition.

As this is a sufficient condition, its violation at this point can be seen as merely incidental to the failure of the scheme. Nevertheless, satisfying the sufficient conditions is all that is necessary to turn an ineffective strategy into a successful one. In our case, if we have quantum control strategies available that are capable to slowing down the evolution of the data qubit to a fraction $\kappa < 1$ of the speed of change of the spectator, so that it now sees an effective error parameter $\kappa\phi$, we find an easier solution to be satisfied:

$$\langle \phi_s^2 \rangle \gg \kappa \left\langle \left| \frac{1}{F^{(2)}(0)} \int_0^{\phi_s} dx (x - \phi_s)^2 F^{(3)}(\kappa x) \right| \right\rangle. \quad (2.17)$$

As long as $F^{(3)}(\phi)$ is continuous near the origin – a feature expected for coherent errors – the right-hand side of the inequality should go to zero as $\kappa \rightarrow 0$. This means we can always find a sufficiently small value of κ so that the sufficient conditions are satisfied. Such a suppression κ of the errors in the data qubit, which can be achieved via control techniques, is equivalent to making the spectator qubit more sensitive to errors. This could be achieved by using different kinds of species of qubits for data and spectators. Although this is a promising direction for future research, in this work we will focus on examples where the control techniques are responsible for the suppression.

Moreover, if the actual error parameter θ is small, we can alternatively suppress ϕ_s by making the data qubit perceive a quadratic error in θ , while the error in the spectator remains linear. In other words, an effective error proportional to $\sim (\theta^2 - \hat{\theta}^2)$ will be smaller than an effective error proportional to $\sim (\theta - \hat{\theta})$.

In the simple example above we can make a feedback scheme work by increasing the relative sensitivity of the spectator to the noise. We will see in the results below that this kind of suppression is also useful in more realistic scenarios. It is

particularly convenient that the quadratic suppression of errors is common in many quantum control techniques [KV09a; SÁS12; Kab+14; Pok+18].

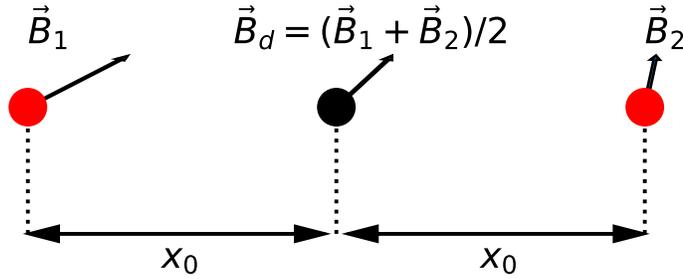
2.2.2 Application to magnetic field noise

A qubit precessing around an axis in the Bloch sphere due to some external coherent error source will behave in a manner that is analogous to a spin-1/2 subjected to an external classical magnetic field. Calling this external classical field \mathbf{B} , the error will be described by the unitary $U(t) = e^{-it\mathbf{B}\cdot\boldsymbol{\sigma}}$, where we are incorporating any constants into the magnitude of \mathbf{B} .

If we know the direction of the classical field \mathbf{B} , we can achieve perfect dynamical decoupling by applying π -pulses in a direction $\hat{\mathbf{n}}$ that is perpendicular to \mathbf{B} [VKL99]. If we do not keep track of the direction of \mathbf{B} , protection against first-order errors can still be obtained via repetitions of an XYXY sequence of π -pulses [GBC90; VKL99], also known as XY-4 [Pok+18; Her+18] or modified CPMG [Mau86]. However, if we acquire information about the direction of \mathbf{B} and rotate the X and Y pulses to a new plane $x'y'$ that is perpendicular to \mathbf{B} , this new tailored $X'Y'$ -4 sequence will not only cancel perfectly the errors caused by a static \mathbf{B} , but will also be robust against small changes in the direction of the classical field.

By placing spectator qubits around the data, as depicted in Fig. 2(a), we can detect drifts in the direction of a magnetic field. We measure the components of $\mathbf{B} = B_x\hat{\mathbf{x}} + B_y\hat{\mathbf{y}} + B_z\hat{\mathbf{z}}$, by suppressing the undesirable parts of the qubit evolution [SSD16] via dynamical decoupling – a process that can be extended to the spectroscopy of non-unitary errors as well [Byl+11] and intermediate situations that involve both kinds of errors [Her+18]. To achieve this, we measure one component at a time, applying π -pulses in the direction that we want to measure, and preparing and measuring the spectator qubit in two distinct bases that are not eigenvectors of the pulses applied. Meanwhile, the data qubit must undergo a dynamical decoupling that suppresses the

(a) Spectator qubits measuring a classical field



(b) Spectator qubits measuring laser beam instability

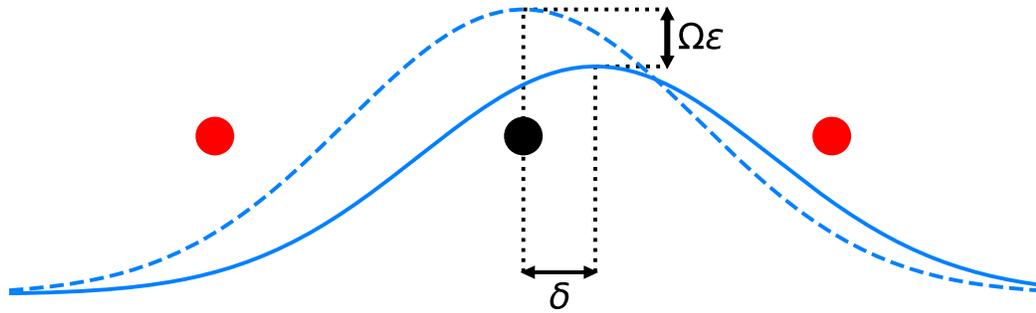


Figure 2.2: **Diagrams of two spectator qubit applications.** In both, we assume an equal distance x_0 between spectators (red) and data qubit (black). In (a), a classical field is assumed to vary linearly in the position coordinate, so the field in the data (\mathbf{B}_d) can be estimated as the average of the field in the equidistant spectators (\mathbf{B}_1 and \mathbf{B}_2). In (b), a laser beam has its ideal Gaussian profile (dashed) changed into an actual beam (solid), which is characterized by the error parameters δ and ϵ .

linear terms of all the components of the magnetic field.

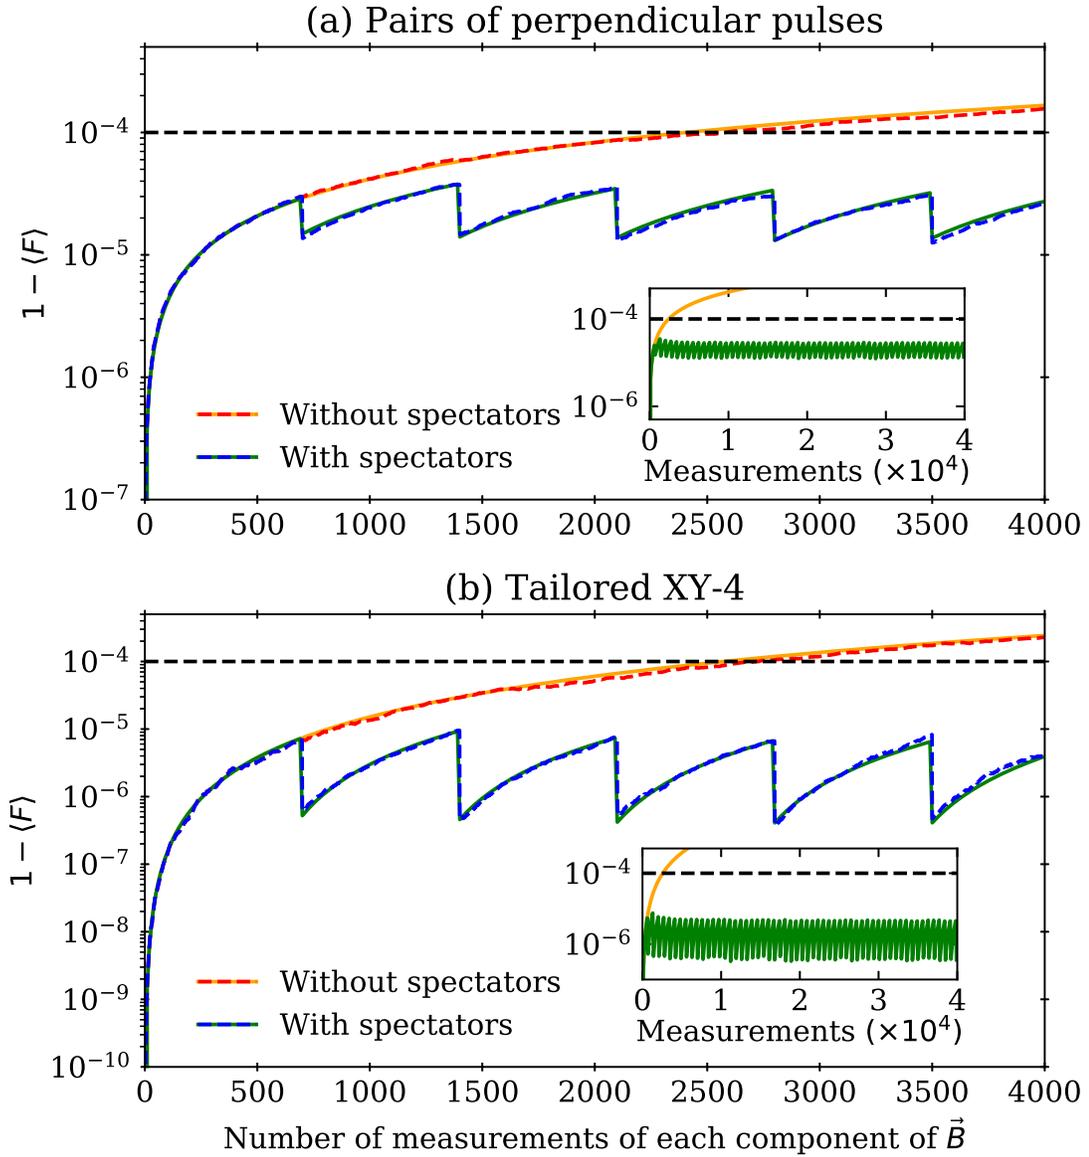


Figure 2.3: **Classical field error with and without spectator qubits.** Here we show the average process fidelity per sequence of four π -pulses spaced by a period τ , calculated numerically and averaged over 1000 runs (dashed), analytically (solid for the case without spectator qubits), and semi-analytically (solid for the case with spectator qubits) when we apply (a) just pulses perpendicular to the direction of the field; (b) a tailored XY-4 sequence where the xy -plane is chosen so that it is perpendicular to the magnetic field. Insets show the long term behavior of the fidelity, where the spectators stay indefinitely below the threshold. We assume the π -pulses to be instantaneous.

In Fig. 3, we compare the process fidelity (2.5) for the case where we maintain the initial calibration with the case where the spectator qubits are used for recalibration. Spectator qubits are able to keep $1 - \langle F \rangle$ below a 10^{-4} threshold after the non-recalibrated system has crossed it. Although some codes have thresholds of the order of 1% [Kni05; RH07; RHG07; FSG09], a more strict threshold would allow fault-tolerance using fewer resources. We consider the dynamical decoupling via sequences of pairs of pulses perpendicular to the direction of the magnetic field and also the tailored X'Y'-4 sequence. In both cases, the spectator qubits stabilize at a level that remains indefinitely below the threshold.

2.2.3 Application to laser beam instability

In ion trap quantum computers, the laser beams used to drive gates, cool ions, and detect states can suffer from common calibration issues such as beam pointing instability and intensity fluctuations [HRB08]. Moreover, they can cause crosstalk, rotations on neighboring qubits that occur when the laser beam overlaps with more qubits than the one being addressed. In principle, the amplitude and pointing instability can be probed by measuring the neighbors [Qia+19], although in practice a series of such measurements can affect other qubits in the chain, creating an additional source of errors which we discuss further in the next section. If we assume the system allows non-disruptive measurements of single qubits, two spectators closely surrounding a data qubit become a possible way of assessing laser beam miscalibrations, as depicted in Fig. 2(b).

Variations in the amplitude of the laser beam change the Rabi frequency Ω by an amount $(1 - \varepsilon)$. Moreover, small errors in the direction of the laser beam are responsible for underrotations. Assuming a Gaussian form for the laser beam, as illustrated in Fig. 2(b), a small pointing displacement of δ results in a quadratic

change in the amplitude Ω of the laser beam affecting the data qubit:

$$\Omega(1 - \varepsilon) e^{-\delta^2} \approx \Omega(1 - \varepsilon) (1 - \delta^2). \quad (2.18)$$

At a distance $\pm x_0$ from the center of the Gaussian, the spectator qubits sense a change in amplitude that is linearly proportional to the pointing displacement δ :

$$\Omega(1 - \varepsilon) e^{-(\pm x_0 - \delta)^2} \approx \frac{\Omega}{c} (1 - \varepsilon) \left(1 \pm 2\delta\sqrt{\ln c}\right), \quad (2.19)$$

where $c = e^{x_0^2}$. This allows the spectator qubits to be sensitive enough to estimate δ before this pointing error grows too much in the data qubit. For ε , the problem of having linear errors both in the data and in the spectator qubits can be overcome by applying composite pulse sequences such as SK1 [MB14]. This kind of sequence reduces the effect of the error in the data qubit to a higher order, while preserving the linear effect on the spectator qubits.

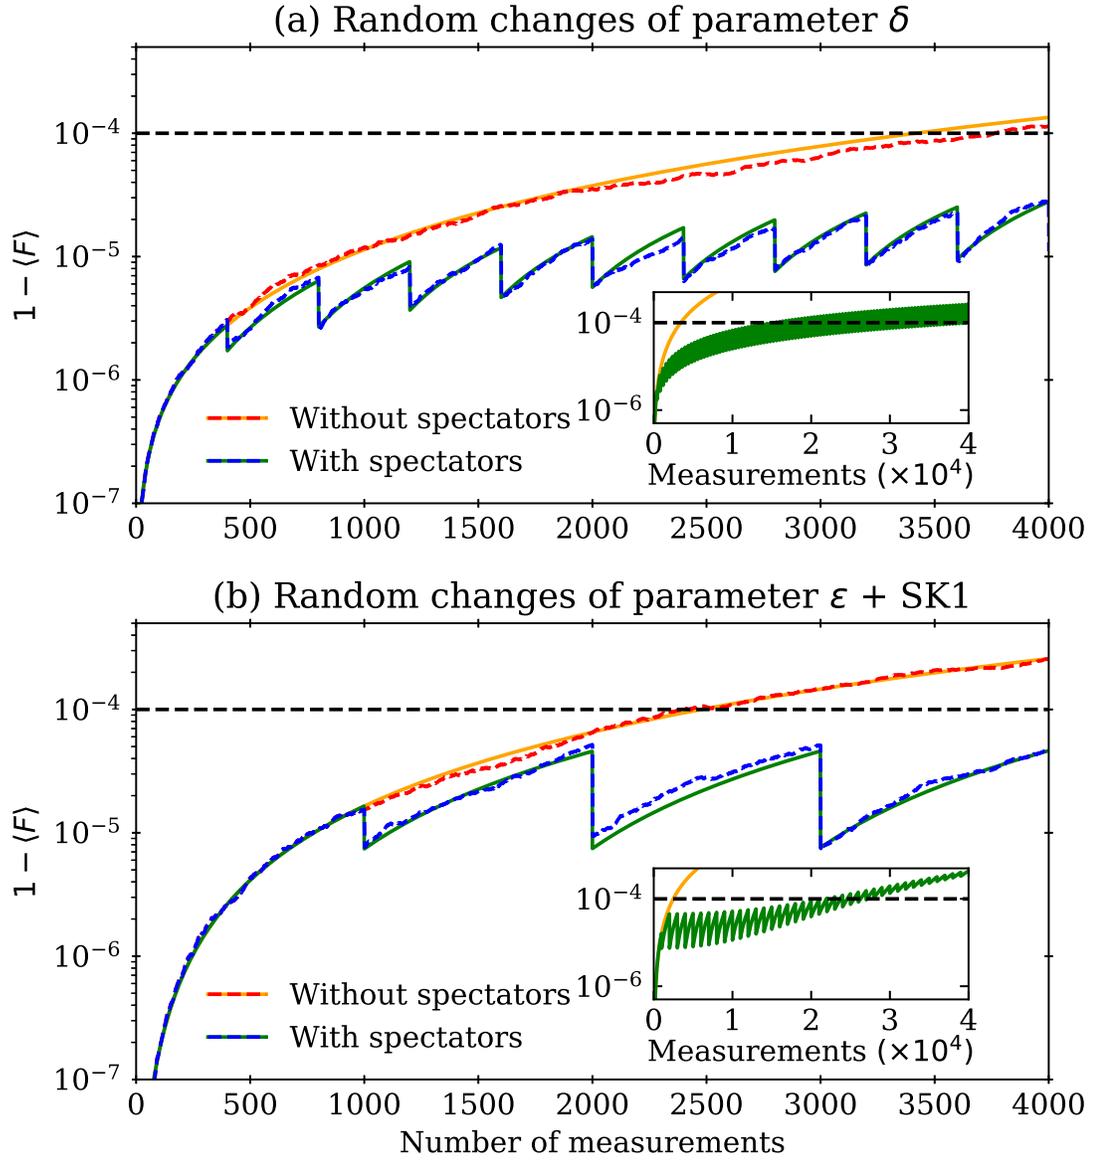


Figure 2.4: **Beam instability errors, with and without spectator qubits.** This is the evolution in time of the average process fidelity per gate for fluctuating parameters (a) δ ; and (b) ε . Discontinuities in the average solution over 1000 numerical runs (dashed) and the analytical approximate solutions (solid) represent points where there is a recalibration. Insets show analytical solutions for the longer time scales, showing the point where the recalibrated systems cross the threshold.

The process fidelity for the laser instability with and without recalibration with spectator qubits is shown in Fig. 4. The improvement in fidelity means that we

are acquiring information fast enough to be able to recalibrate the system before the errors become too large. If spectator qubits and data both were subjected to an error linearly proportional to ε , the recalibration would not be able to keep the errors under the same threshold for the same values of the parameters, as can be seen in Fig. 5. It is therefore crucial to choose a measurement strategy that balances the rate of acquisition of information and the rate with which the errors increase. We define τ_{nospec} to be the time when $\langle F_{\text{nospec}} \rangle$ crosses this threshold and τ_{spec} when $\langle F_{\text{spec}} \rangle$ crosses the threshold. In Fig. 6, we show which combinations of random walk parameters and measurements per spectator cycle M are still capable of providing an effective recalibration mechanism. Fig. 6 also functions as a control landscape and a map that shows the frequency of updates that is adequate for a given order of magnitude of the error. Using our estimate of the order of magnitude of the error obtained from the initial calibration, we can find the adequate frequency of update of the control method by looking for the blue regions below the solid black line.

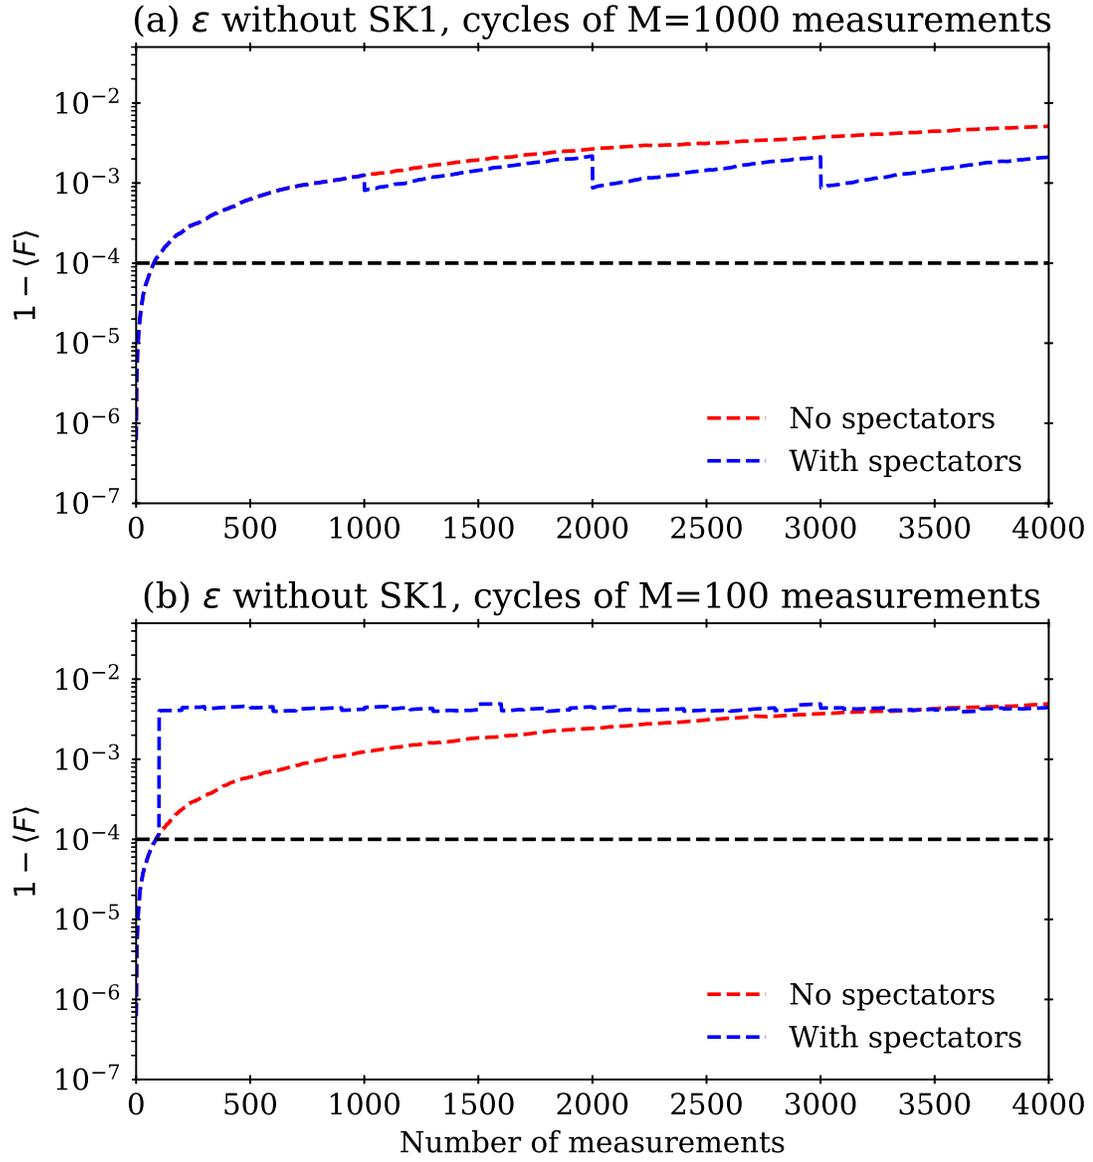


Figure 2.5: **Example of a frustrated spectator qubit scheme.** This numerical simulation is averaged over 1000 runs for the same situation as Fig. 4 (b), but with the SK1 turned off, so that the effect of the error parameter ε is linear in both spectators and data qubits. Under these circumstances, it is never possible to keep the average error under the 10^{-4} threshold: either (a) we take too long to use the information from the spectator qubits and the correlation is already lost, or (b) we update before sufficient Fisher information is available, causing further miscalibration of the system.

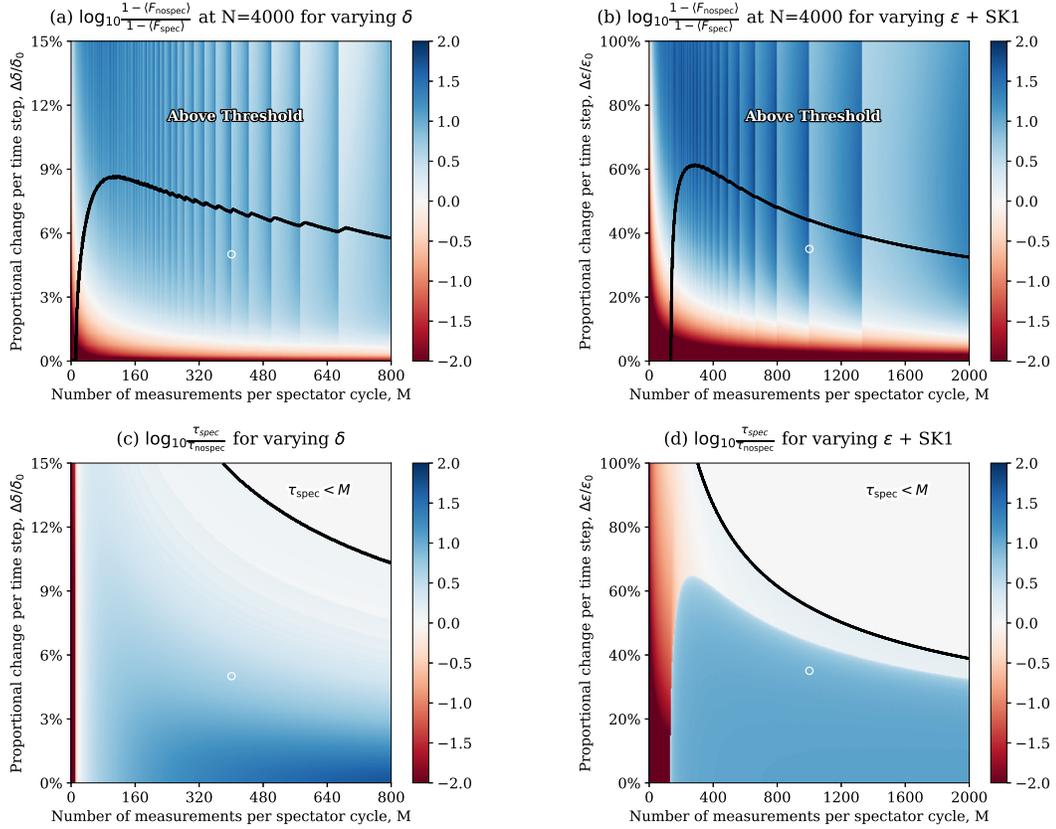


Figure 2.6: **Control landscapes for beam instability errors.** These are heat maps in base-10 logarithm of the: (a,b) ratio of $1 - \langle F \rangle$ with and without spectator qubits at fixed time $N = 4000$; and (c,d) ratio between the points where the 10^{-4} threshold is crossed, with and without spectator qubits. These control landscapes are heat maps showing the effect of spectator qubits for different values of measurements per cycle (M) and step size. Initial values of the error parameters and ion distance are the same as in Fig. 4, and the values of $\Delta\epsilon$, $\Delta\delta$, and M that correspond to those from Fig. 4 are marked by the white circle. Blue regions represent settings where the spectators improve either (a,b) the fidelity; or (c,d) time it stays below the 10^{-4} threshold. Regions above the black curve represent situations where: (a,b) the system with spectators cross the 10^{-4} threshold for the process fidelity; or (c,d) where the 10^{-4} threshold is crossed before there has been time to complete the first spectator cycle. Spectator qubits perform worse (red areas) when very few measurements are performed before updating (left extremity of the graphs), or when the rate of change is so small that not recalibrating is a better strategy, as in the bottom of (a) and (b). Discontinuities along the x -axis in (a) and (b) correspond to situations where the end of a spectator cycle occurs at the point $N = 4000$, and are analogous to the discontinuities seen in Fig. 4.

2.3 Discussion

We have shown that spectator qubits are capable of recalibrating an error reduction scheme for coherent errors with a precision that is only limited by the Fisher information available and by our capacity of slowing down the rate with which the error in the data qubit changes with time. In this work, both spectator and data qubits were equally sensitive to noise, and we had to use quantum control methods to slow down the error accumulation on data qubit. One benefit of such a strategy is that we can dynamically allocate spectator qubits [Gup+19], and use all physical qubits in our system as spectator or data qubits. But we can sacrifice this advantage to construct (or choose) spectator qubits that are more sensitive to a particular type of noise than our data qubits. Such a strategy eliminates the need for quantum control methods that were used to slow down errors on data qubit. One example is to use multi-species ions in ion-trap architecture [BB19]. We can use Zeeman qubits (first-order magnetic field sensitive) as our spectator qubits and hyperfine qubits (first-order magnetic field insensitive) as our data qubits for dealing with magnetic field noise.

In the case of the spectator qubits used to calibrate dynamical decoupling, in Fig. 3, we have seen the fidelity remain under the threshold for an indefinite period. We believe this is possible because the error in this setting depends mainly on the angle between the classical field and the pulses, a parameter whose value is not allowed to grow indefinitely.

For laser beam instability, the insets in Fig. 4 show that even when we are using spectator qubits, the average gate fidelity crosses over the threshold at a later time. We believe this is because the error parameters become very large as the random walk is unbounded. This contrasts with the magnetic field parameters, whose random walk was bounded. When the error parameters become very large ($\varepsilon, \delta > 1$), the

data qubit (error is quadratic) becomes more sensitive to the error than the spectator qubits (error is linear). One possible way to fix this is to include an external classical controller that restricts the maximum variance of the fluctuating error parameters and prevents the crossing of the threshold.

It is worth noting for the laser beam instability case that although we have simulated the fidelity of a single gate (σ_x) due to miscalibration, it is straightforward to extend our approach to an arbitrary computation. We can do this by interleaving cycles composed of gates that we want to calibrate on data qubits and spectator qubit measurements between gates of the algorithm.

When attempting to put the beam instability feedback loop into practice, it may become important to take into account the additional errors that can be caused by measuring neighboring qubits. However, in this setting where we attempt to improve the process fidelity per gate, the penalty for the measurement errors does not accumulate, rather impacting the performance of the gates individually. If a measurement error is modeled by an incoherent process that occurs with probability p , this will count as an incoherent channel that occurs after the gate and reduces the probability of no error occurring by a fraction $1 - p$. This new background error will reduce overall gate fidelity and in order to maintain gates below the threshold error, it is critical to maintain an even better calibration. In effect this will lower the threshold that we have to reach by a fixed amount, but the overall analysis of the problem remains the same. When the measurement error occurs during an SK1 sequence, the interaction between the series of gates and the incoherent channel will be more complicated, but the overall effect will still be that an incoherent error spoils the state of the system with probability p , having the overall effect of lowering the threshold.

The possibilities of applications of spectator qubits are not limited to the two coherent errors that we simulated above. Protection against magnetic fields, for ex-

ample, besides being relevant to ion traps and nuclear spin qubits, could be extended to detection and dynamical decoupling of a classical external electric field \mathbf{E} for qubits that are instead sensitive to electric fields, such as antimony nuclei [Asa+19].

In future full-fledged quantum computing systems, spectator qubits will enable error rates below the threshold for fault-tolerance for longer times than systems without spectator recalibration. This will allow for longer quantum computations. For near and medium-term applications, however, enhancements would be required in order to reduce the prohibitive number of measurements necessary to obtain a reliable estimate of the change in the calibration. It would be particularly desirable to implement small corrections in the calibration after fewer measurements, possibly assuming some prior knowledge of how the calibration changes, or a specific biased drift of the error parameters. These could be combined with other venues for improvement, such as using Bayesian learning protocols [GB18; Gup+19] to make more accurate previsions of future evolution of error parameters or to implement adaptive measurements [Gra+12; DCS17], and using entangled states [Eld+18], many-body Hamiltonians [DCS17], quantum codes [Zho+18], or optimal control [LY17] to maximize the information available.

2.4 Methods

2.4.1 Error model

In our numerical, analytical and semi-analytical simulations for Figs. 3, 4, 5, and 6, the error parameters θ are assumed to start at a fixed value θ_0 and fluctuate in time according to a random walk with unbiased Gaussian steps of average size $\Delta\theta$, so that the probability of it having a value θ_N after N steps will be:

$$p(\Theta = \theta | \Theta_0 = \theta_0) = \mathcal{N}(\theta_0, N(\Delta\theta)^2; \theta), \quad (2.20)$$

where the random variable Θ_n gives the value of the error parameter after n steps, and $\mathcal{N}(\mu, \sigma^2; x)$ is the normal distribution with mean μ and variance σ^2 :

$$\mathcal{N}(\mu, \sigma^2; x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left\{ -\frac{(x - \mu)^2}{2\sigma^2} \right\}. \quad (2.21)$$

Suppose that, given an actual value of the error parameter θ and an estimate ϑ , we know the expression of the process fidelity per gate, $F(\theta, \vartheta)$. Then, if the parameter drifts in time but our estimate is not updated, we can use the probability distribution of the random walk to find the average fidelity per gate after N steps when spectators are not recalibrating the system, F_{nospec} :

$$\langle F_{\text{nospec}} \rangle = \int d\theta p(\Theta_N = \theta | \Theta_0 = \theta_0) F(\theta, \vartheta). \quad (2.22)$$

This expression can be analytically calculated for all the application above (see Supplementary Section III).

If spectator qubits are present, the estimate ϑ is updated after every cycle of M measurements. After the k th cycle of measurements of the spectators, the next estimate is obtained via an estimator $\theta_{kM}^{(\text{est})}$ that consists of the average of the error parameter sampled at the previous M steps of the random walk:

$$\Theta_{kM}^{(\text{est})} = \frac{\Theta_{(k-1)M+1} + \Theta_{(k-1)M+2} + \dots + \Theta_{kM}}{M}, \quad (2.23)$$

where we assume that the parameters change sufficiently slowly so that θ has a precisely defined value during each measurement. For this reason, the variance of the Gaussian in Eq. (2.20) can always be rescaled so that the number of steps of the random walk matches the number of measurements.

The probability distribution of the estimator will be a Gaussian, as this is a random variable consisting of the average of the Gaussian random variables Θ_{kM+n} :

$$p(\Theta_{kM}^{(\text{est})} = \theta_{kM}^{(\text{est})} | \Theta_{kM} = \theta_{kM}) = \mathcal{N}(\mu_k, \sigma_k^2; \theta_{kM}^{(\text{est})}). \quad (2.24)$$

Therefore, the probability distribution will be entirely characterized by the two cumulants that can be calculated from Eq. (2.23), which are (see Supplementary Section II) the mean μ_k :

$$\mu_k = \theta_{kM} + \frac{M-1}{kM} \frac{\theta_0 - \theta_{kM}}{2}, \quad (2.25)$$

and the variance σ_k^2 :

$$\sigma_k^2 = \frac{M-1}{3} \frac{4kM - 3M - 3k + 2}{4kM} (\Delta\theta)^2. \quad (2.26)$$

Finally, the measured value $\bar{\theta}$ may differ from the estimator, according to the lower limit of the Cramér–Rao bound (2.2), by an amount that corresponds to the inverse of the Fisher information f_θ times the number of measurements:

$$p(\bar{\Theta} = \bar{\theta} | \Theta_{kM}^{(\text{est})} = \theta_{kM}^{(\text{est})}) = \mathcal{N}(\theta_{kM}^{(\text{est})}, (Mf_\theta)^{-1}; \bar{\theta}). \quad (2.27)$$

Given these probability distributions for θ_{kM} , $\theta_{kM}^{(\text{est})}$, and $\bar{\theta}$, the average fidelity N steps after the k th spectator cycle, which we call F_{spec} , can be calculated from the average fidelity for a fixed calibration given in Eq. (2.22):

$$\begin{aligned} \langle F_{\text{spec}} \rangle = & \int d\theta_{kM} \left\{ p(\Theta_{kM} = \theta_{kM} | \Theta_0 = \theta_0) \right. \\ & \times \int d\theta_{kM}^{(\text{est})} \left[p(\Theta_{kM}^{(\text{est})} = \theta_{kM}^{(\text{est})} | \Theta_{kM} = \theta_{kM}) \right. \\ & \left. \left. \times \int d\bar{\theta} p(\bar{\Theta} = \bar{\theta} | \Theta_{kM}^{(\text{est})} = \theta_{kM}^{(\text{est})}) \langle F_N(\theta_0, \bar{\theta}) \rangle_0 \right] \right\}. \quad (2.28) \end{aligned}$$

Using the assumption that the error parameters are small, we solved the triple integrals analytically for ε and δ (see Supplementary Section IV) and numerically for the magnetic field case.

2.4.2 Magnetic field noise

In the simulation of the dynamical decoupling of a magnetic field, we assumed the field gradient to be linear, so that measurements in two spectators are sufficient to determine the field in the data qubit. We choose the $\hat{\mathbf{z}}$ axis to coincide with the initial direction of the magnetic field. Using τ to denote the time spacing between instantaneous π -pulses, we choose the initial value \mathbf{B}_1 of the magnetic field in one of the spectators to satisfy $\tau\mathbf{B}_1 = 2 \cdot 10^{-3} \hat{\mathbf{z}}$ when the data qubit undergoes a 2-pulse sequence, and $\tau\mathbf{B}_1 = 3.8 \cdot 10^{-2} \hat{\mathbf{z}}$ when it undergoes the tailored XY-4 sequence. The second spectator is assumed to experience initially half of the value of this magnetic field ($\mathbf{B}_2 = \mathbf{B}_1/2$).

Each component of the magnetic field was assumed to perform an independent random walk, with steps of different size. We choose standard deviations $\Delta B_x/B_{1,x} = 3\%$, $\Delta B_y/B_{1,y} = 2\%$, $\Delta B_z/B_{1,z} = 1\%$ for each random walk. These components are then assessed separately and sequentially in the spectator qubits, which is done by preparing and measuring the spectator in eigenbases of two distinct Pauli matrices that are perpendicular to the component of \mathbf{B} that we want to measure. There is no need for spectator qubit re-initialization after each measurement as measurement in a particular Pauli basis prepares the spectator qubit in an eigenbases of that Pauli operator. The other components of \mathbf{B} are decoupled by applying a sequence of π -pulses to the spectators between each measurement, so that we can approximate our estimates of B_x , B_y , and B_z by:

$$\tilde{B}_x = -\arcsin\left(\frac{\langle 0|U^\dagger(n\tau)\sigma_y U(n\tau)|0\rangle}{4n\tau}\right), \quad (2.29)$$

$$\tilde{B}_y = \arcsin\left(\frac{\langle 0|U^\dagger(n\tau)\sigma_x U(n\tau)|0\rangle}{4n\tau}\right), \quad (2.30)$$

$$\tilde{B}_z = \arcsin\left(\frac{\langle +|U^\dagger(n\tau)\sigma_y U(n\tau)|+\rangle}{4n\tau}\right), \quad (2.31)$$

where $\langle \psi | U^\dagger(n\tau) \sigma_i U(n\tau) | \psi \rangle$ represents the averages of a measurement of σ_i in a system prepared at a state $|\psi\rangle$ and left to evolve for a time $n\tau$. The number n of π -pulses before each measurement was chosen as 20 for spectators aiding the perpendicular 2-pulse sequence, and 4 for spectators whose information was used to tailor a XY-4 sequence. After $M = 700$ measurement cycles, we use the new estimate of the direction of \mathbf{B} to update our dynamical decoupling control parameters. For the pairs of perpendicular π -pulses, we make the pulse direction perpendicular to \mathbf{B} . We estimate the plane that is normal to \mathbf{B} , find two perpendicular pulses in that plane, and use them as our tailored XY-4 sequence. We repeat this process throughout the length of our computation.

2.4.3 Laser beam instability

For spectator qubits used to reduce the underrotation caused by pointing instability of the laser beam, we simulate a series of σ_x gates applied to the data qubit and assume the presence of either a fluctuating parameter δ or ε . The δ parameter is assumed to start the random walk at $\delta_0 = 0.02$ and proceed with Gaussian steps of standard deviation $\Delta\delta/\delta_0 = 5\%$, while the ε starts at $\varepsilon_0 = 0.002$ and proceeds with steps of standard deviation $\Delta\varepsilon/\varepsilon_0 = 35\%$. We assume an initial calibration that allows us to estimate δ to a precision $\bar{\delta}/\delta_0 = 99\%$, and ε to a precision $\bar{\varepsilon}/\varepsilon_0 = 75\%$.

The δ errors naturally cause a greater effect in the spectators than in the data, as can be seen from Eqs. (2.18) and (2.19), where the difference is between a linear and a quadratic dependence. After four σ_x gates, we measure the spectator qubits, which we assume to be at a distance $x_0 = (\ln 12)^{1/2}$ from the data (x_0^2 is measured in units of twice the variance). To maximize the Fisher information for uncorrelated probes, the two spectator qubits are prepared and measured in an eigenstate of σ_z . The averages $\langle \sigma_z \rangle_1, \langle \sigma_z \rangle_2$ of the M measurement results in spectator qubits 1 and 2

are then used to estimate δ according to:

$$\bar{\delta} = \frac{1}{2x_0} \frac{\arccos(\langle\sigma_z\rangle_1) - \arccos(\langle\sigma_z\rangle_2)}{\arccos(\langle\sigma_z\rangle_1) + \arccos(\langle\sigma_z\rangle_2)}, \quad (2.32)$$

which follows from Eq. (2.19). As only δ^2 affects the data, the sign of our estimate of δ is irrelevant. After $M = 400$ repeated measurements, we build sufficient confidence in our estimate $\bar{\delta}$ so that, for future gates, we adjust the Rabi frequency to $\Omega/(1 - \bar{\delta}^2)$ using our classical control setup to compensate for the pointing instability.

As the parameter ε is linear in all qubits, we apply an SK1 composite pulse sequence [BHC04] to slow down the error accumulation in the data qubit. Measurements on the spectator qubits – assumed to be at a distance $x_0 = (\ln 1.8)^{1/2}$ from the data (where x_0^2 is measured in units of twice the variance) – are performed after each regular σ_x gate is applied, but before the application of the second and third pulses of the SK1 sequence. The value of ε is then estimated from the measurement results of $\langle\sigma_z\rangle_1, \langle\sigma_z\rangle_2$:

$$\bar{\varepsilon} = 1 - e^{x_0^2} \frac{\arccos(\langle\sigma_z\rangle_1) + \arccos(\langle\sigma_z\rangle_2)}{2\pi/(1 - \tilde{\varepsilon})}, \quad (2.33)$$

where $\tilde{\varepsilon}$ is the previous estimate of ε . After $M = 1000$ measurements, we update the Rabi frequency to $\Omega/(1 - \bar{\varepsilon})$ to compensate for the errors.

2.5 Key Observations

We have shown that spectator qubits are capable of recalibrating an error reduction scheme for coherent errors with a precision that is only limited by the Fisher information available and by our capacity of slowing down the rate with which the error in the data qubit changes with time. In the case of the spectator qubits used to calibrate dynamical decoupling, in Fig. 3, we have seen the fidelity remain under the threshold for an indefinite period. We believe this is possible because the error in

this setting depends mainly on the angle between the classical field and the pulses, a parameter whose value is not allowed to grow indefinitely.

For laser beam instability, the insets in Fig. 4 show that even when we are using spectator qubits, the average gate fidelity crosses over the threshold at a later time. We believe this is because the error parameters become very large as the random walk is unbounded. This contrasts with the magnetic field parameters, whose random walk was bounded. When the error parameters become very large ($\varepsilon, \delta > 1$), data qubit (error is quadratic) becomes more sensitive to the error than the spectator qubits (error is linear). One possible way to fix this is to include an external classical controller that restricts the maximum variance of the fluctuating error parameters and prevents the crossing of the threshold.

It is worth noting for the laser beam instability case that although we have simulated the fidelity of a single gate (σ_x) due to miscalibration, it is straightforward to extend our approach to an arbitrary computation. We can do this by interleaving cycles composed of gates that we want to calibrate on data qubits and spectator qubit measurements between gates of the algorithm.

The possibilities of applications of spectator qubits are not limited to the two coherent errors that we simulated above. Protection against magnetic fields, for example, besides being relevant to ion traps and nuclear spin qubits, could be extended to detection and dynamical decoupling of a classical external electric field \mathbf{E} for qubits that are instead sensitive to electric fields, such as antimony nuclei [Asa+19].

In future full-fledged quantum computing systems, spectator qubits will enable error rates below the threshold for fault-tolerance for longer times than systems without spectator recalibration. This will allow for longer quantum computations. For near and medium-term applications, however, enhancements would be required in order to reduce the prohibitive number of measurements necessary to obtain a reliable estimate of the change in the calibration. It would be particularly desirable

to implement small corrections in the calibration after fewer measurements, possibly assuming some prior knowledge of how the calibration changes, or a specific biased drift of the error parameters. One way to approach this is to start with single shot measurement and correction which is what we discuss in the next section. These could be combined with other venues for improvement, such as using Bayesian learning protocols [GB18; Gup+19] to make more accurate previsions of future evolution of error parameters or to implement adaptive measurements [Gra+12; DCS17], and using entangled states [Eld+18], many-body Hamiltonians [DCS17], quantum codes [Zho+18], or optimal control [LY17] to maximize the information available.

2.6 Single shot correction theory

In this section, we propose an alternative protocol that is more appropriate for near-term quantum systems. To do so, we first introduce a probabilistic noise model and present an optimization routine that can be utilized to determine the optimal single-shot corrections for this noise. Next, we outline a practical end-to-end protocol to estimate these corrections. To support our findings, we present detailed numerical simulations and offer insights into the differences between our original protocol and the single-shot correction protocol.

Let $U(\theta)$ be the unitary that describes the coherent error on the system, where θ is a random variable with probability density $p(\theta)$. We assume the density $p(\theta)$ and the dependence of $U(\theta)$ to θ are known, but the actual value of θ is unknown. To obtain information about θ we perform measurements on the spectator qubits. Suppose these measurements yield outcome $m \in M$, with probability $p(m|\theta)$, where M is a finite set of possible outcomes (For instance, it could be a finite bit string corresponding to the outcomes of the measurements on spectator qubits). The conditional distribution $p(m|\theta)$ is determined by our choice of state preparation and measurement on the spectator qubits. Then, for each possible outcome m , we choose an appropriate

unitary V_m , which approximately cancels the effect of $U(\theta)$, such that $V(m)U(\theta)$ is approximately equal to the identity operator for most values of θ and m . To summarize, our protocol for correcting coherent errors is determined by a conditional probability $p(m|\theta)$ and the choice of unitaries $\{V_m\}$. To quantify the performance of a protocol we can use the average gate fidelity or, equivalently, entanglement fidelity. Recall that the entanglement fidelity for a channel \mathcal{E} with Kraus operators $\{E_k\}$, is given by

$$F_e(\mathcal{E}) = \langle \Psi | [\mathcal{E} \otimes \mathcal{I}] (|\Psi\rangle\langle\Psi|) | \Psi \rangle = \frac{1}{d^2} \sum_k |\text{Tr}(E_k)|^2 \quad (2.34)$$

where $|\Psi\rangle$ is a maximally-entangled state, and d is the dimension of the system Hilbert space. From here we can calculate the average gate fidelity

$$\begin{aligned} \int d\psi \langle \psi | \mathcal{E} (|\psi\rangle\langle\psi|) | \psi \rangle &= \frac{d + \sum_k |\text{Tr}(E_k)|^2}{d(d+1)} \\ &= \frac{1 + d \times F_e(\mathcal{E})}{d+1}, \end{aligned} \quad (2.35)$$

where $d\Psi$ is the uniform (unitarily invariant) measure over pure states. To quantify the overall error after applying the unitary correction V_m , we consider the average entanglement fidelity. Assuming the conditional probability $p(m|\theta)$ is given and fixed, we write this average as

$$\begin{aligned} \bar{F}(\{V_m\}) &= \int d\theta p(\theta) \sum_{m \in \mathcal{M}} p(m|\theta) \times F_e(V_m U_\theta) \\ &= \int d\theta \sum_{m \in \mathcal{M}} p(\theta, m) \times \frac{1}{d^2} |\text{Tr}(V_m U_\theta)|^2 \end{aligned} \quad (2.36)$$

where $F_e(V_m U_\theta)$ denotes the entanglement fidelity for the unitary transformation $V_m U_\theta$, and $p(\theta, m) = p(m|\theta)p(\theta)$. The fact that there are a finite number of spectator qubits, limits the form of the conditional distribution $p(m|\theta)$ that can be achieved by

performing measurements on these qubits. Our final goal is to optimize over all such realizable conditional distributions $p(m|\theta)$ and unitaries $\{V_m\}$ to find the maximum achievable value of average entanglement fidelity

$$\max_{p(m|\theta)} \max_{\{V_m\}} \bar{F}(\{V_m\}) \quad (2.37)$$

$$= \max_{p(m|\theta)} \max_{\{V_m\}} \int d\theta p(\theta) \sum_{m \in \mathcal{M}} p(m|\theta) \times \frac{1}{d^2} |\text{Tr}(V_m U_\theta)|^2 \quad (2.38)$$

This optimization problem can be solved in two steps: first, we assume $p(m|\theta)$ is fixed and given, and find the set $\{V_m\}$ which maximizes the average gate fidelity for this $p(m|\theta)$, and then we maximize over the set of all possible conditionals $p(m|\theta)$. In the first step of the optimization problem, where $p(m|\theta)$ is fixed, we need to solve

$$\bar{F}_{\text{opt}} = \max_{\{V_m\}} \int d\theta p(\theta) \sum_{m \in \mathcal{M}} p(m|\theta) \times \frac{1}{d^2} |\text{Tr}(V_m U_\theta)|^2 \quad (2.39)$$

$$\text{subject to : } V_m V_m^\dagger = I, \forall m \in \mathcal{M} \quad (2.40)$$

It is useful to note that we can optimize over unitaries V_m independently, i.e. for each m we can find the optimal

$$\max_{V_m} \int d\theta p(\theta) p(m|\theta) \times \frac{1}{d^2} |\text{Tr}(V_m U_\theta)|^2 \quad (2.41)$$

$$\text{subject to : } V_m V_m^\dagger = I \quad (2.42)$$

and the set of these unitaries will form the optimal set $\{V_m\}$.

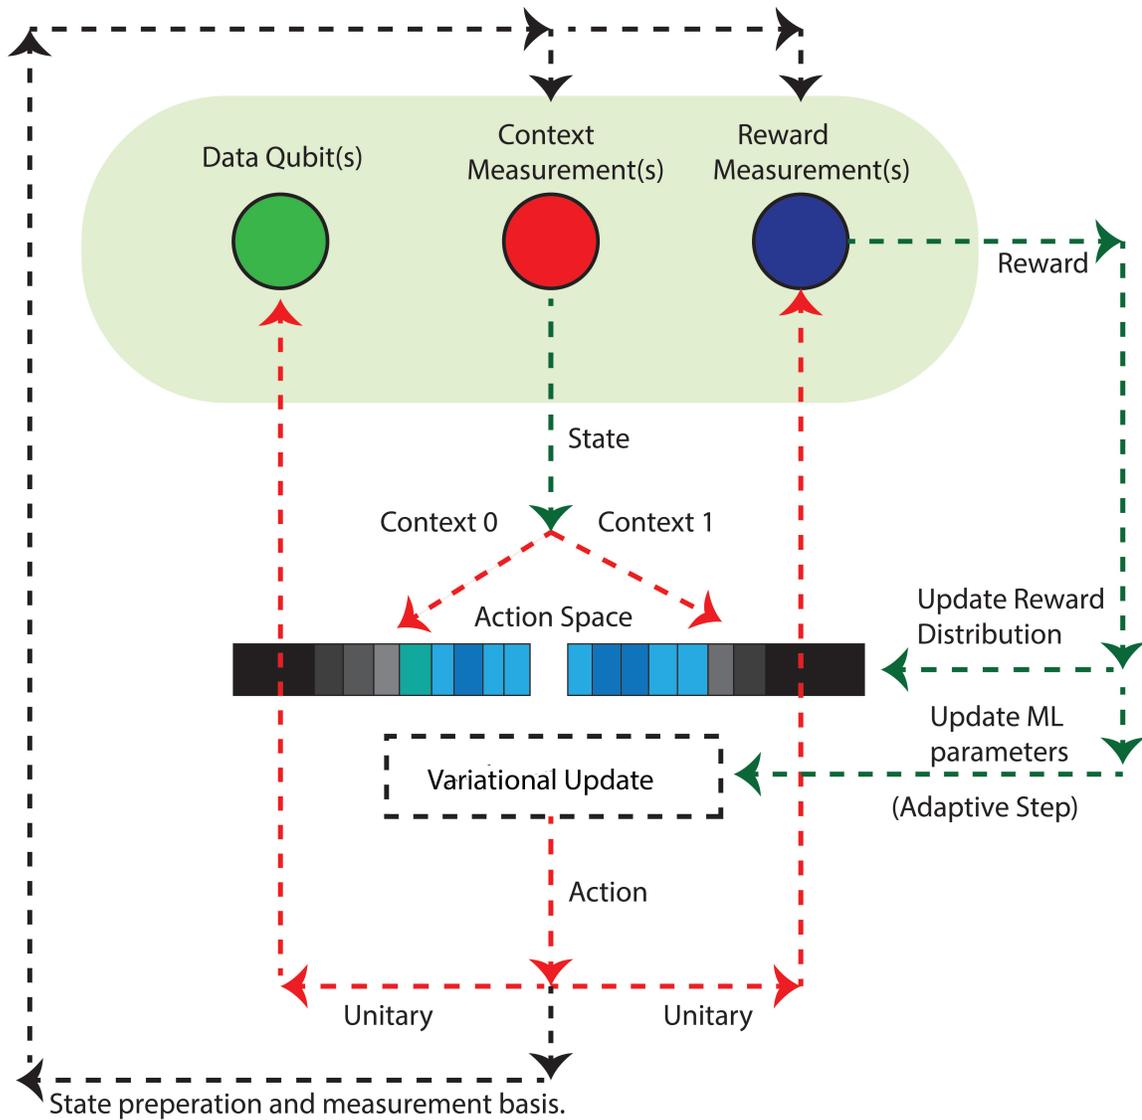


Figure 2.7: **Single shot spectator qubit protocol** Here we have two "types" of spectator qubits. One that determines the optimal measurement and one that determines the optimal correction. The optimal measurement and correction are parameterized and learned using a variational algorithm.

2.7 Low overhead spectator qubit algorithm

The strategy above for mitigating coherent errors come with its own set of challenges when we apply this in a real system. Practically, it is hard to estimate the unknown error probability distribution. Further, what if the distribution changes? We need

a mechanism to detect such changes. Once we have detected changes, we have to re-run optimizations again. This technique also assumes the correcting unitaries are not noisy. We need a separate technique to learn or mitigate errors in those corrections. While the theoretical solution gives us an optimization problem to solve in order to find the optimal measurement and the set of correcting unitaries, it requires us to have complete information about the underlying error channel (error direction and error distribution). We have designed a spectator qubit based protocol, as illustrated in Fig 2.7, that does not require this prior information but instead calculates the optimal measurement choice and the set of correcting unitaries as sub-routines.

Algorithm 1 describes this algorithm. We essentially have three types of spectator qubits to (1) determine the optimal measurement (these spectators are called context determining spectators denoted by N_S and the optimal measurement is parameterized by θ_S), (2) determine the correcting set of unitaries (these spectators are called reward spectator denoted by N_R and the optimal measurement is parameterized by $\{\theta_R^i\}$ - two unitaries because we are dealing with two outcome measurements), and (3) perform the optimal measurement (called N_C). Optimal θ_S and θ_R are calculated using parameterized unitaries and optimized using two loss functions \mathcal{L}_S and \mathcal{L}_R . Hyperparameters of this algorithm are gradient step sizes (η_R and η_S) and batch sizes (τ_R and τ_S) need to be tuned. Intuitively, \mathcal{L}_S is chosen to maximize the variance of the context measurements, measured using N_S and \mathcal{L}_R is chosen to maximize the average gate fidelity measured using N_R .

2.8 Results

For numerical simulations, here we consider a special case where the noise is just a random rotation around a fixed, but arbitrary (known) axis. We assume the noise unitary is of the form

$$U_\theta = e^{i\sigma_3\theta} \tag{2.43}$$

Algorithm 1 Variational spectrators (with majority polling) as diagrammed in Fig 2.8

- 1: Given total spectator qubit count N , allocate spectrators between N_S , N_R and N_C .
 - 2: Initialize parameters $\{\theta_R^j\}_{j=1}^2$ and θ_S . Choose step-sizes η_R and η_S , and batch sizes τ_R and τ_S .
 - 3: Select tangent space $|+\rangle$ which maximizes discrimination information as a prior.
 - 4: Environment selects classical noise process.
 - 5: **for all** time-steps t **do**
 - 6: Prepare N_C and N_S spectrators in $|+\rangle$.
 - 7: Prepare N_R spectrators in Haar-random $|\psi\rangle$.
 - 8: Environment draws a random error parameter and applies the corresponding single-qubit unitary error to each spectrator and data qubit.
 - 9: Sample $\langle L_{\theta_S} \rangle_{N_C}$ using $|+\rangle$.
 - 10: **if** $\langle L_{\theta_S} \rangle_{N_C} < 0.5$ **then**
 - 11: Apply mitigating unitary $U(\theta_R^1)$ to data.
 - 12: **else**
 - 13: Apply $U(\theta_R^2)$ to data.
 - 14: **end if**
 - 15: **if** $t \bmod \tau_R = 0$ **then**
 - 16: $\theta_R^j \leftarrow \theta_R^j - \eta_R \langle \nabla_{\theta_R^j} \mathcal{L}_R(\theta_R^j) \rangle_{N_R \tau_R}$ using $|\psi\rangle$.
 - 17: **end if**
 - 18: **if** $t \bmod \tau_S = 0$ **then**
 - 19: $\theta_S \leftarrow \theta_S - \eta_S \langle \nabla_{\theta_S} \mathcal{L}_S(\theta_S) \rangle_{N_S \tau_S}$ using $|+\rangle$.
 - 20: **end if**
 - 21: **end for**
-

computation and the second protocol in improving quantum memory.

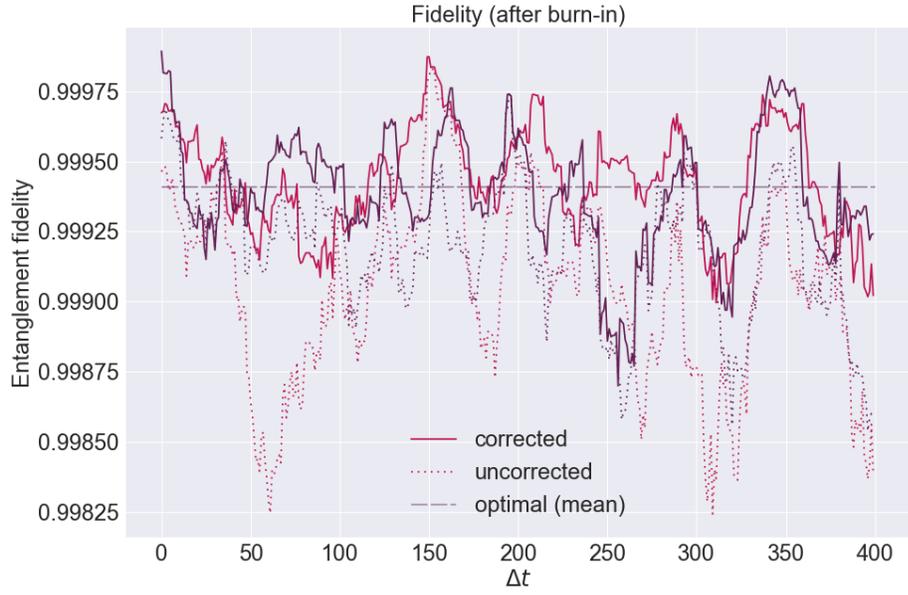


Figure 2.9: **Stationary Regime** Performance of our proposed algorithm in the high fidelity regime. Here the error samples are drawn from a fixed distribution $\mathcal{N}(0, \pi/32)$. Fidelities are plotted after the initial calibration phase. Different colors represent different error realizations.

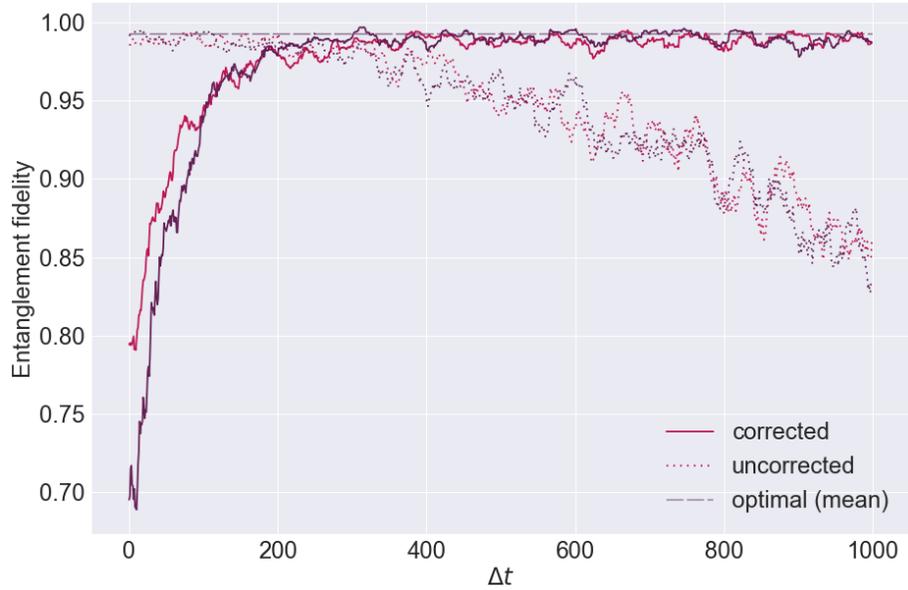


Figure 2.10: **Non-stationary Regime** Performance of our proposed algorithm in the non-stationary regime. Here the error samples are drawn from a distribution that is changing over time. We find the spectator qubit protocol is able to adapt as the noise distribution drifts. Different colors represent different error realizations.

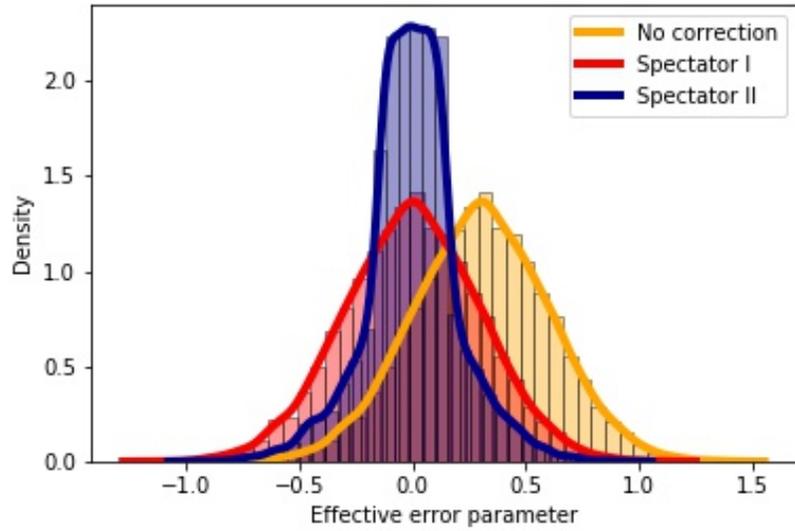


Figure 2.11: **Effective error parameter** Comparing the probability density function of the effective error parameters between (1) no correction (orange) (2) first spectator qubit protocol (red) and (3) second spectator qubit protocol (blue)

Circuit level error mitigation with hidden inverses

The majority of this chapter is taken from Ref [Maj+22; LMP22]:

Swarnadeep Majumder, Christopher G Yale, Titus D Morris, et al. “Characterizing and mitigating coherent errors in a trapped ion quantum processor using hidden inverses”. arXiv preprint arXiv:2205.14225

and

Vicente Leyton-Ortega, **Swarnadeep Majumder**, and Raphael C Pooser. “Quantum error mitigation by hidden inverses protocol in superconducting quantum devices”. Quantum Science and Technology.

3.1 Introduction

Small algorithms on quantum computing testbed devices can help diagnose performance problems when the resulting data is viewed within the context of dominant low-level noise sources. In near-term pre-fault-tolerant devices, coherent noise (such as over- or under-rotation errors) can lead to an array of incorrect results at the algorithmic level, including incorrect expectation values in time evolution simulations or slow convergence and incorrect parameter determination in variational algorithms. Here, we use prototypical variational quantum chemistry calculations as a diagnostic algorithm to probe coherent noise on a trapped ion testbed platform. Two methods of noise mitigation - randomized compiling [WE16] and the application of hidden inverse gates [Zha+22] - are used to both characterize and mitigate time-dependent

coherent noise in test circuits. We present simulations of one and two-qubit noise models on trapped ion quantum devices, validated by experimental data, to verify this hypothesis.

Quantum errors in a quantum computer can be thought of as any evolution of the qubits that differs from the ideal intended operation. There are different ways one may categorize quantum errors, but for the purpose of this work we limit our discussion to two distinct types of errors: incoherent error and coherent error. Incoherent errors map pure states to mixed states resulting in loss of purity often caused by classical noise. Some well known quantum channels such as depolarization and dephasing are examples of incoherent errors. Conversely, coherent error can be represented as a single unitary operator and maps pure states to pure states. Such an error usually comes from miscalibration or drift out of calibration of the control system used to drive the qubit operations. If one has a good understanding of the system interaction Hamiltonian, it is often possible to describe coherent errors as additive or multiplicative terms in the control parameters. This allows one to use targeted characterization procedures to learn about the error parameters which in turn can be used for better calibration. While fault tolerant protocols using a quantum error correction (QEC) code have been demonstrated recently [Ega+21; Kri+21; Rya+21], noisy intermediate scale quantum (NISQ) devices are unable to take advantage of QEC because of high noise levels and small system sizes. For NISQ devices, one can try to mitigate these errors to minimize their impact on the quantum computer’s output; such protocols are collectively known as EM (error mitigation) protocols.

Multiple quantum characterization, verification, and validation (QCVV) protocols exist for characterizing the noise channels and error models discussed here, including gate set tomography [Blu+13], randomized benchmarking [Joh+15], generalized model fitting [Nie+21], and many others. While these frameworks are very

general, typically allowing for deduction of a diverse array of noise models, a large quantity of experimental data can be required to learn the models in some cases. The exclusive access to many quantum computing platforms required to obtain the data is not always readily available. Recent applications of various EM techniques in the NISQ era provide a promising middle ground approach [Nat+21; Kim+21; PLP21; Str+21; Piv+21; LaR+20; Zha+20; Cza+21; Suz+22]. Essentially, the efficacy of an EM routine is an indicator of the presence of a specific type of noise present in the machine, and with sufficient *a priori* information these techniques can be used to characterize noise in a quantum processor.

3.2 Hidden inverses protocol to mitigate coherent errors

Hidden inverses (HI) [Zha+22; LMP22; Yet+21] is an EM protocol that addresses coherent errors via noise cancellation. HI does not require any additional gates or any post processing. The key idea behind HI is that a few of the common unitaries used in quantum computing are self-adjoints. These self-adjoint unitary operators are constructed from primitive gates which themselves are not self-adjoints. As an example if we want to apply a certain gate $G = ABC$, we can either apply $G = ABC$ or $G^\dagger = C^\dagger B^\dagger A^\dagger$. While in the absence of noise, G and G^\dagger implement the same physical operation; this is not true in general when A , B , C , A^\dagger , B^\dagger and C^\dagger are subjected to errors. Given the freedom to apply G or G^\dagger , a compiler needs to choose which one to apply based on other gates in the vicinity within a quantum circuit along with an understating of the noise process. Developing compiler logic for picking the optimal configuration given an arbitrary circuit and a noise model is an open problem and one that we will not discuss in this research. Instead, we focus our attention on demonstrating hidden inverses’s performance with hand-crafted circuits.

Hadamard (H) and controlled-NOT (CNOT) are two self-inverse unitary operations widely used in quantum computing. We can decompose their standard or native

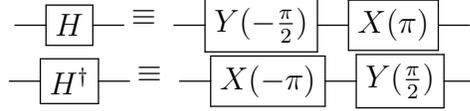


Figure 3.1: Standard and Hermitian conjugated decompositions of H gate with native trapped ion quantum operations consisting of single qubit gates

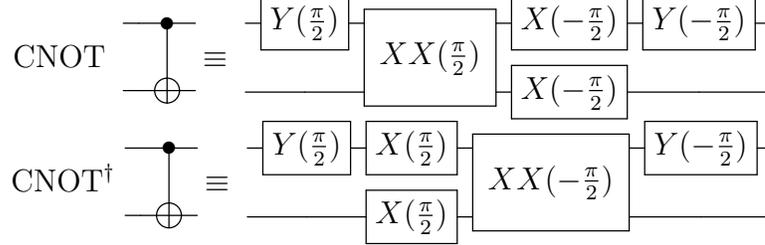


Figure 3.2: Standard and Hermitian conjugated decompositions of CNOT gate with native trapped ion quantum operations consisting of single qubit gates and Mølmer-Sørensen interactions.

configurations into trapped-ion native single qubit gates and Mølmer-Sørensen (MS) interactions (Fig 3.1, 3.2).

3.3 Experimental implementation of quantum circuits on a trapped ion processor

We investigate these hidden inverse protocols on the Quantum Scientific Computing Open User Testbed (QSCOUT), a room-temperature quantum processor based on trapped ions [Cla+21] housed at Sandia National Laboratories. This investigation consisted of either one or two qubits, in which the qubit states comprise the hyperfine ‘clock’ transition of a $^{171}\text{Yb}^+$ ion, $^2\text{S}_{1/2} |F=0, m_F = 0\rangle (|0\rangle)$ and $|F=1, m_F = 0\rangle (|1\rangle)$ [Olm+07]. The ions are trapped in a radio-frequency (RF) pseudopotential generated on a microfabricated surface electrode trap, the Sandia High Optical Access (HOA2.1) trap [Mau16], with radial frequencies, $\omega_{r,i}/2\pi = 2.2$ and 2.5 MHz, and an axial frequency, $\omega_a/2\pi = 700$ kHz. This axial frequency sets the spacing of the ions, $4.5 \mu\text{m}$, to match the spacing of the individual addressing beams. Each ion is then imaged in a separate core of a multicore fiber array via magnification optics to match

the fiber core spacing to the 4.5 μm separation. The ions are cooled and detected along their $^2\text{S}_{1/2}$ to $^2\text{P}_{1/2}$ transition using 370 nm light.

All gates performed in the system are based on Raman transitions from a pulsed 355 nm laser [Hay+10] with beams in a counter-propagating configuration, consisting of a wide ‘global’ beam illuminating all ions and individual-ion addressing beams generated via a multi-channel acousto-optic modulator (AOM) [Deb+16] counter to the global beam. Both the multi-channel AOM and the single-channel AOM for the global beam are driven by a custom coherent control system, Octet, providing two tones for each AOM channel.

Gates are fully parameterized within the QSCOUT system, providing both arbitrary phase and rotation angle for single- and two-qubit gates. Single-qubit gates are driven using two Raman tones on the appropriate individual beam in a co-propagating Raman configuration. Our two-qubit MS gate is generated with two tones on each participating ion’s individual beam and another tone on the global beam. The three tones are required to form the appropriate Raman transitions symmetrically detuned from a red and blue motional sideband. The MS gate as implemented in QSCOUT additionally has a series of single-qubit basis transformations to eliminate phase instabilities between the counter- vs. co-propagating configurations. The bare $\text{MS}(\theta)$ gate is an XX interaction, $\text{XX}(\theta) = e^{-i\frac{\theta}{2}\sigma_x \otimes \sigma_x}$, that exists in a counter-propagating basis. By surrounding the two-qubit gate with the appropriate counter-propagating single-qubit gates, we transform the XX interaction into a ZZ interaction to eliminate phase instabilities with the co-propagating single-qubit gates [Lee+05]. The gate is then further surrounded with co-propagating single-qubit gates to complete the transformation back into an XX interaction within the co-propagating basis.

The physical single-qubit gates’ (R_X and R_Y) amplitudes are square-pulsed and gapless, meaning there is no “off” period between single-qubit gates. Their rotation

angle is determined by the duration of the pulse. R_Z gates are virtual, treated as cumulative phase shifts by Octet.

The MS gate amplitude is Gaussian-shaped, and the angle of rotation is controlled by the global beam amplitude while accounting for distortions and saturation in the amplifier and AOM. A MS gate with a negative rotation angle is achieved by flipping the phase on one of the two qubits by π radians. For all gates, phases are set by the relative phase difference between two of the RF tones applied to the ion. The MS gate also incurs an AC Stark shift which is cancelled through a dynamically evolving virtual Z rotation applied throughout the duration of the MS gate pulse. This ability to fully control both the phase and rotation angle for single- and two-qubit gates is needed for these investigations to intentionally introduce static coherent rotation and phase errors. Typical physical single-qubit gate fidelities for a rotation angle of $\pi/2$ are estimated to be $99.5\pm 0.3\%$, while two-qubit $XX(\pi/2)$ are $97\pm 1\%$.

Additionally, to cool the ions to near the motional ground state, the ions are both Doppler cooled and resolved-sideband cooled [Des+04]. Typically, 60 loops of cooling are performed on all relevant sidebands to reach the minimum motional state; however, the number of cooling loops can be reduced so the ion will begin the desired circuit in a higher motional mode, inducing additional errors. These errors include an overall static coherent under-rotation error and shot-to-shot coherent errors. At very small numbers of cooling loops ($\lesssim 10$ loops), there is an increase in the population of non-desirable qubit states - i.e. for a state beginning in $|00\rangle$, an MS gate ideally creates a state somewhere within the $|00\rangle$ and $|11\rangle$ basis, but instead there are nonzero populations in the $|10\rangle$ and $|01\rangle$ states. These unwanted populations can be as large as $\sim 5\%$ when no cooling loops are applied.

The gate sequences are all programmed using Just Another Quantum Assembly Language (Jaqal), a flexible programming language designed to support the underlying hardware of the QSCOUT system [Mor+20]. Since Jaqal contains the fully

parameterized gateset, we can introduce intentional coherent errors directly at the Jaqal level as different rotation angles or phases, rather than underlying waveforms, described by JaqalPaw (Pulses and Waveforms), the pulse-level counterpart to Jaqal [Lob+21]. On the other hand, errors introduced by insufficient cooling of the ion are instituted by adapting the cooling cycle in the waveform language JaqalPaw.

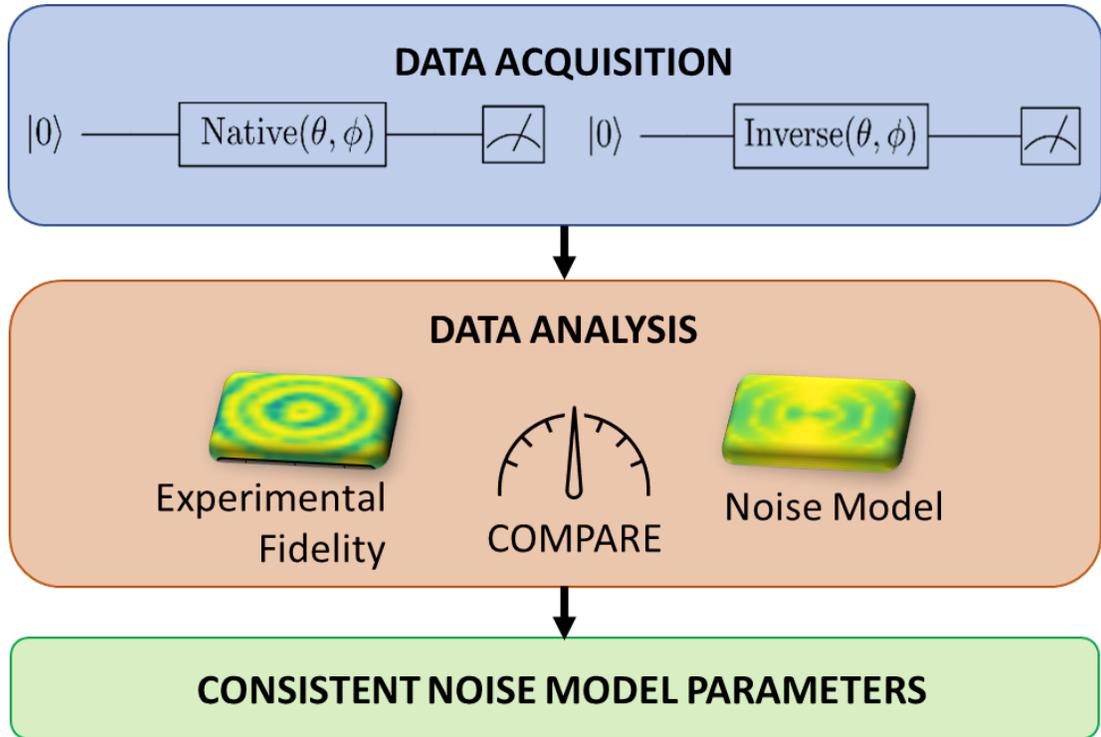


Figure 3.3: Illustration of noise characterization using single qubit hidden inverses.

3.4 Demonstration of hidden inverses as an error characterization experiment with single qubit gates

3.4.1 Experimental design

We use H , H^\dagger , and a few single qubit rotations to develop an error characterization experiment. Here we summarize the key insight of the protocol. If we compare outputs of two circuits (1) $H - H$ (native circuit) and (2) $H - H^\dagger$ (inverted circuit) under coherent noise model that inverts with the inverse of the gate, we find circuit

(2) will result in error cancellation while circuit (1) will amplify the error. In order to investigate different sources of coherent errors in the system, we inject parameterized single qubit rotations in between the gates as follows (1) $H - X(\Theta) - Z(\Phi) - H$, $H - Y(\Theta) - Z(\Phi) - H$ or (2) $H - X(\Theta) - Z(\Phi) - H^\dagger$, $H - Y(\Theta) - Z(\Phi) - H^\dagger$. When we sweep over $\{\Theta, \Phi\}$, we will amplify or cancel various types of coherent errors in the native and inverted circuits. For some choices of $\{\Theta, \Phi\}$ the inverse circuits will outperform the native circuits, while for others the native circuits will be a better choice. Aided with numerical simulations, patterns in the fidelity (or population) plots of the various circuits in phase space of $\{\Theta, \Phi\}$ will reveal the underlying noise process.

There are two stages in this experiment (i) data collection stage and (ii) data analysis stage. In the data collection stage, the circuits were run 100 times before a measurement, i.e. (1) $[H - X(\Theta) - Z(\Phi) - H, H - Y(\Theta) - Z(\Phi) - H]^{100}$ or (2) $[H - X(\Theta) - Z(\Phi) - H^\dagger, H - Y(\Theta) - Z(\Phi) - H^\dagger]^{100}$ to amplify any effects of cancellation. 21 equally spaced points were chosen in between $-\frac{\pi}{36} \leq (\Theta, \Phi) \leq \frac{\pi}{36}$ for a total of 441 points in a 2D-grid for each circuit. Each data set was taken with random ordering in Θ and ϕ to average out effect of any long term drift. Outcomes were averaged over 200 shots taken back-to-back for a given (Θ, Φ) . This data was taken without dynamical decoupling or composite pulses.

In order to characterize system parameter drift and verify the model, $[H - X(\Theta) - Z(\Phi) - H^\dagger]^{100}$ was also run multiple times back to back for four different cases (1) natural drift: no re-calibration was performed between experiments (2) re-calibration: re-calibration was performed between experiments (3) intentional amplitude noise injection (4) intentional phase noise injection.

$$\begin{aligned}
\text{NativeXZ} &\equiv \{ \boxed{H} \boxed{X(\Theta)} \boxed{Z(\Phi)} \boxed{H} \}^{100} \\
\text{InverseXZ} &\equiv \{ \boxed{H} \boxed{X(\Theta)} \boxed{Z(\Phi)} \boxed{H^\dagger} \}^{100} \\
\text{NativeYZ} &\equiv \{ \boxed{H} \boxed{Y(\Theta)} \boxed{Z(\Phi)} \boxed{H} \}^{100} \\
\text{InverseYZ} &\equiv \{ \boxed{H} \boxed{Y(\Theta)} \boxed{Z(\Phi)} \boxed{H^\dagger} \}^{100}
\end{aligned}$$

Figure 3.4: **Standard and Hermitian conjugated circuits for single qubit error characterization.** $\{\}^{100}$ denotes the circuit is repeated 100 times before measurement.

3.4.2 Classical simulation and noise model

Outcomes from stage (1) were first converted into fidelity. To calculate fidelity, we performed noisy and ideal simulations of the quantum circuits. For the noisy simulation, we assume a physics-focused, single-qubit, ion-trap noise model parameterized by over-rotation error (ϵ), phase error (ϕ), and detuning error (δ). In this model, ideal $X_{ideal}(\theta = \Omega t) = e^{-i\frac{\Omega t}{2}X}$ and $Y_{ideal}(\theta = \Omega t) = e^{-i\frac{\Omega t}{2}Y}$ where Ω is the Rabi frequency, become:

$$X_{noisy} = e^{-i\frac{\Omega(1+\epsilon)t}{2}(Cos(\phi)X+Sin(\phi)Y)+\frac{\delta t}{2}Z}$$

$$Y_{noisy} = e^{-i\frac{\Omega(1+\epsilon)t}{2}(Cos(\phi)Y+Sin(\phi)X)+\frac{\delta t}{2}Z}.$$

3.4.3 Result

We found the best noise parameters by curve fitting the experimental data with noisy simulation using SciPy's non-linear least-square fitting algorithm [Vir+20]. We used $\{\Theta, \Phi\}$ as the two-dimensional independent variables and the population as the dependent variable. Fig 3.5 displays the population phase space for $[H - X(\Theta) - Z(\Phi) - H^\dagger]^{100}$ circuit. In order to characterize drifting of error parameters in time, we plot the estimated error parameters as a function of number of runs in Fig 3.6. We find the variance of the parameter estimate from the curve-fit to be in the order of 10^{-4} , hence they are omitted in the plots. Long term drifts in ϵ while ϕ and

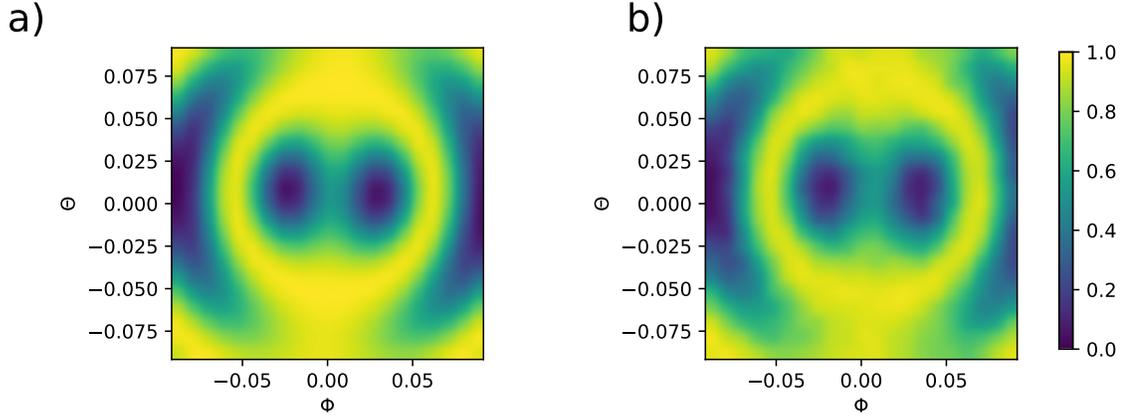


Figure 3.5: **Phase space population (probability of measuring $|0\rangle$) after applying $[H - X(\Theta) - Z(\Phi) - H^\dagger]^{100}$ on $|0\rangle$ as a function of $\{\Theta, \Phi\}$ for (a) simulation and (b) experiment**

δ approximately stays constant. Interleaved re-calibration routine largely stabilizes ϵ . With the addition of intentional amplitude and phase noise injection over many runs, we find the protocol is able to accurately track the error parameters.

3.5 Hidden inverses as an error mitigation protocol for applications: variational quantum eigensolvers

3.5.1 Hamiltonian and Ansatz

Variational quantum eigensolvers (VQE) are a successful cornerstone of hybrid quantum-classical algorithms explored on NISQ era machines. For a given system Hamiltonian H , and parameterized ansatz $|\Psi(\alpha)\rangle$, the ground state and its energy are obtained by classically minimizing the energy,

$$E(\alpha) = \langle \Psi(\alpha) | H | \Psi(\alpha) \rangle . \quad (3.1)$$

Thus they form an ideal platform to explore the efficacy of mitigation algorithms like hidden inverse and randomized compiling. In order to pursue this goal, we explore the VQE in the minimal basis for equilibrium molecular hydrogen. We use the Brayvi-Kitaev mapping, along with qubit tapering to yield a spin Hamiltonian

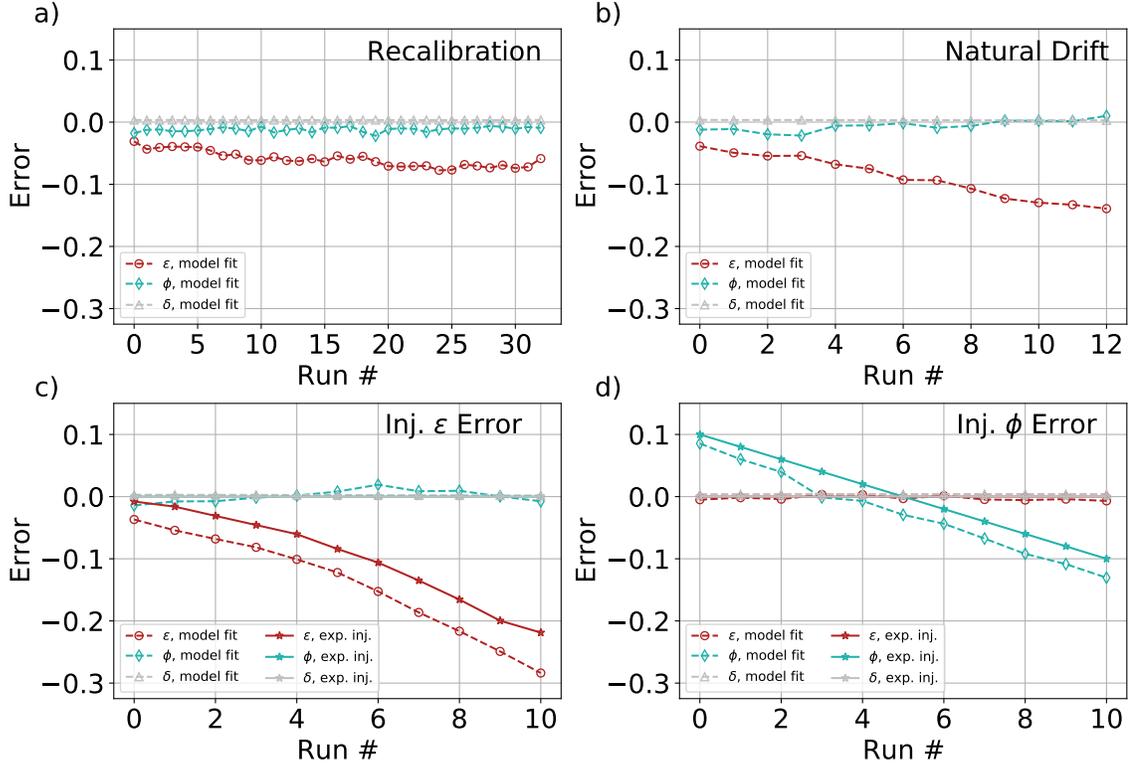


Figure 3.6: **Error parameter drift** Error parameters as a function of number of runs are plotted. (a) is for the case where system parameters were calibrated in between runs and we find the estimated error parameters stable. (b) is for natural drift where re-calibration steps were skipped. We find significant drift in the estimated over-rotation (red) parameter. We then inject artificial noise in the form of amplitude error (c) and phase error (d). We find good agreement between the injected noise (dotted line) and estimated parameters.

of the following form,

$$\begin{aligned}
 H = & 0.304794 * II + 0.3555426 * IZ - .485486 * ZI \\
 & + 0.581232 * ZZ + 0.089500 * (XX + YY), \quad (3.2)
 \end{aligned}$$

indicating that energy evaluation requires at most three circuit evaluations for ZZ, XX, and YY expectation values. We test the hidden inverse in this system in a prototypical one parameter two-qubit circuit ansatz in Fig 3.7.

This circuit is easily decomposed into native ion gates as in Fig 3.2, where the default compilation uses the same construction for both CNOTs, and the HI imple-

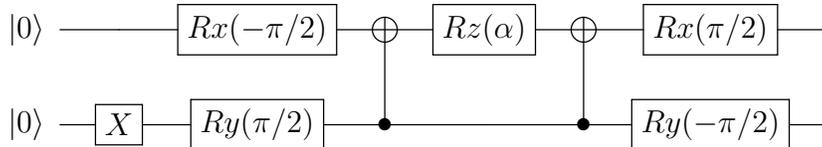


Figure 3.7: **VQE tapered ansatz**

mentation uses the CNOT^\dagger construction for the second CNOT. For the default and HI circuit, we swept over all angles $|\alpha| < \frac{\pi}{2}$, with 200 shots for each evaluation.

3.5.2 Randomized Compiling

We also wish to compare the benefit of HI inverse methods in Fig. 3.7 with other methods of coherent error mitigation, namely randomized compiling. Randomized compiling is a powerful tool, where coherent errors of hard gates, typically those of the two-qubit gates, are suppressed by averaging over unitarily equivalent circuits of the two-qubit gates via application of easy “twirling” gates. In [Zha+22], we simulated performance of randomized compiling and compared with hidden inverses for a set of noise models. We found hidden inverses can improve overall circuit fidelity in certain noise processes. In order to demonstrate how this fidelity improvement translates to improvement in algorithmic performance experimentally, we compare hidden inverses with randomized compiling for our VQE problem. For application to our problem, we identified the CNOT as our hard gates, and the Pauli operators as our twirling group, yielding a total of 256 unique circuit evaluations. We found that randomly choosing 10 circuits from this grouping was sufficient for converged results. For each of the random circuits, they were repeated 20 times to yield the same number of shots to compare with the default and HI circuits.

3.5.3 Purification

To further explore mitigation strategies in conjunction with the circuit level strategies already presented, we utilized fermionic density matrix purification first demon-

strated in quantum computing in [McC+19a]. Here, instead of directly evaluating the energy based on measured Pauli expectations, we note that the 2-dimensional reduced fermionic density matrix, ρ can also be established from the already measured expectation values.

$$\begin{aligned}\rho &= \begin{pmatrix} \langle 01 | \Psi \rangle \langle \Psi | 01 \rangle & \langle 01 | \Psi \rangle \langle \Psi | 10 \rangle \\ \langle 10 | \Psi \rangle \langle \Psi | 01 \rangle & \langle 10 | \Psi \rangle \langle \Psi | 10 \rangle \end{pmatrix} \\ &= \begin{pmatrix} (1 - \langle IZ \rangle)/2 & (\langle XX \rangle + \langle YY \rangle)/4 \\ (\langle XX \rangle + \langle YY \rangle)/4 & (1 - \langle ZI \rangle)/2 \end{pmatrix}. \end{aligned} \quad (3.3)$$

Once measured, this fermionic density matrix contains the effect of noise, and does not represent a pure state. As an effective one-particle pure state, its eigenvalues should be either zero or one. To project this onto a pure state, we diagonalize ρ , identifying the eigenvector $|\phi\rangle$ corresponding to the largest eigenvalue, and then establish a new fermionic density matrix $\rho_p = |\phi\rangle \langle \phi|$. This procedure is not scalable for N-fermion problems, but similarly themed attempt have been explored within enforcing N-representability constraints on measured 2-reduced density matrices [RBM18]. Thus purified energies can then be evaluated using ρ_p .

3.5.4 Noise Injection

To investigate the noise mitigation properties of these diverse approaches, we intentionally introduce errors onto the two-qubit gates in a variety of manners. While the two-qubit gate in the circuit is the CNOT and its inverse, we inject noise on the native gate of the system, the $XX(\pm\pi/2)$ MS gate. As described in section 3.3, to introduce an under- or over-rotation error, we program that error at the circuit level, resulting in a change in the amplitude of the global beam power delivered. We specifically introduce the rotation error symmetrically, so in the CNOT, we utilize $XX(\pi/2 + \epsilon)$ and for the $CNOT^\dagger$, $XX(-\pi/2 - \epsilon)$. This is most akin to the types of rotation errors that would occur naturally in the system, which could be due either

to a miscalibration of the overall power necessary, or a drift in the overall system detuning also causing systematic over- or under-rotation. In Fig. 4.1, we introduce a systematic under-rotation of 0.5 radians (right) in comparison to the nominally fully calibrated approach (left).

Additionally, we can introduce a broader array of errors through reduction in sideband cooling loops as described in section 3.3. In Fig. 3.9, we compare full cooling of 60 loops (left) to a reduction in sideband cooling of 12 loops (right). For full cooling, 60 loops, we see the desired rotation and a residual population in $|01\rangle$ and $|10\rangle$ of $\sim 1.5\%$. However, with only 12 cooling loops applied, we see a static under-rotation of ~ 0.3 radians and population in the $|01\rangle$ and $|10\rangle$ states of $\sim 2\%$. Additionally, we expect to see an increase in the stochastic coherent error with a partially cooled ion.

We performed detailed numerical simulation of the loss landscape with noise injection (solid lines in Fig. 4.1 and 3.9). Our error model consist of ideal single qubit gates, no state prep or measurement errors and all the errors are attributed to the MS gate. Reducing the number of cooling loops affects the effective Rabi frequency and also amplifies the effect of residual entanglement between spin and motion. We model the effective Rabi-frequency change as a combination of static offset to Rabi frequency and a stochastic reduction in Rabi frequency following Debye-Waller effect [Win+98]. This can be viewed as a shot to shot coherent error. On the other hand, we model residual entanglement as two-qubit depolarizing noise on the spin states. For the noiseless case, our best fit model has average phonon number of 0.05 (motional mode heating), a static coherent over-rotation of 0.09 radians (laser intensity miscalibration) and a two-qubit depolarizing probability of 0.02 (residual entanglement between spin and motion) resulting in a 97.5% MS gate fidelity. Next, for coherent rotation noise injection, we find a static effective under-rotation of 0.45 from the ideal gate while the two other noise parameters stay the same resulting in a

91% MS gate fidelity. Finally, for the case where we limit the number of cooling loops to 10-12, we find the best model has average phonon number of 0.5, a static over-rotation of 0.12 radians and a two qubit depolarizing probability of 0.06 resulting in a 89% MS gate fidelity.

3.6 Key observations

The experimental results and corresponding numerical simulations highlight a few key ideas. First, both hidden inverse gates and randomized compiling are effective methods for coherent noise mitigation. They are markedly different in their application; while randomly compiled circuits add extra twirling gates, the use of HI gates preserves the original circuit depth. They have applicability in different scenarios; for example, an inverse CNOT preceded by a CNOT and a large rotation will not provide the desired noise cancellation. In the work outlined here, we showed that in certain applications with small rotational parameters, such as some instances of VQE, the HI method is very effective near the variational minimum. Secondly, noise amplification allows one to test both the efficacy of a noise mitigation technique and also verify hypotheses about the nature of noise present. Here, noise amplification was used to test the capacity of randomized compiling, HI gates, and density matrix purification to mitigate over and under-rotation errors as well as cooling cycle reduction. These error types can all potentially be attributed to a miscalibrated apparatus (or stale calibrations), which may occur in hardware that has extended up time in between calibration cycles. Third, simulations at the single qubit level were used to fit effective noise model parameters to the experimental data, allowing us to verify hypotheses about the source of the noise as well as the behavior of the noise model under HI circuits. In the experimental device, when calibration cycles were explicitly skipped, the model fits predicted that under-rotation errors would be significant source of coherent noise (see Fig. 3.6). Explicitly adding under-rotation

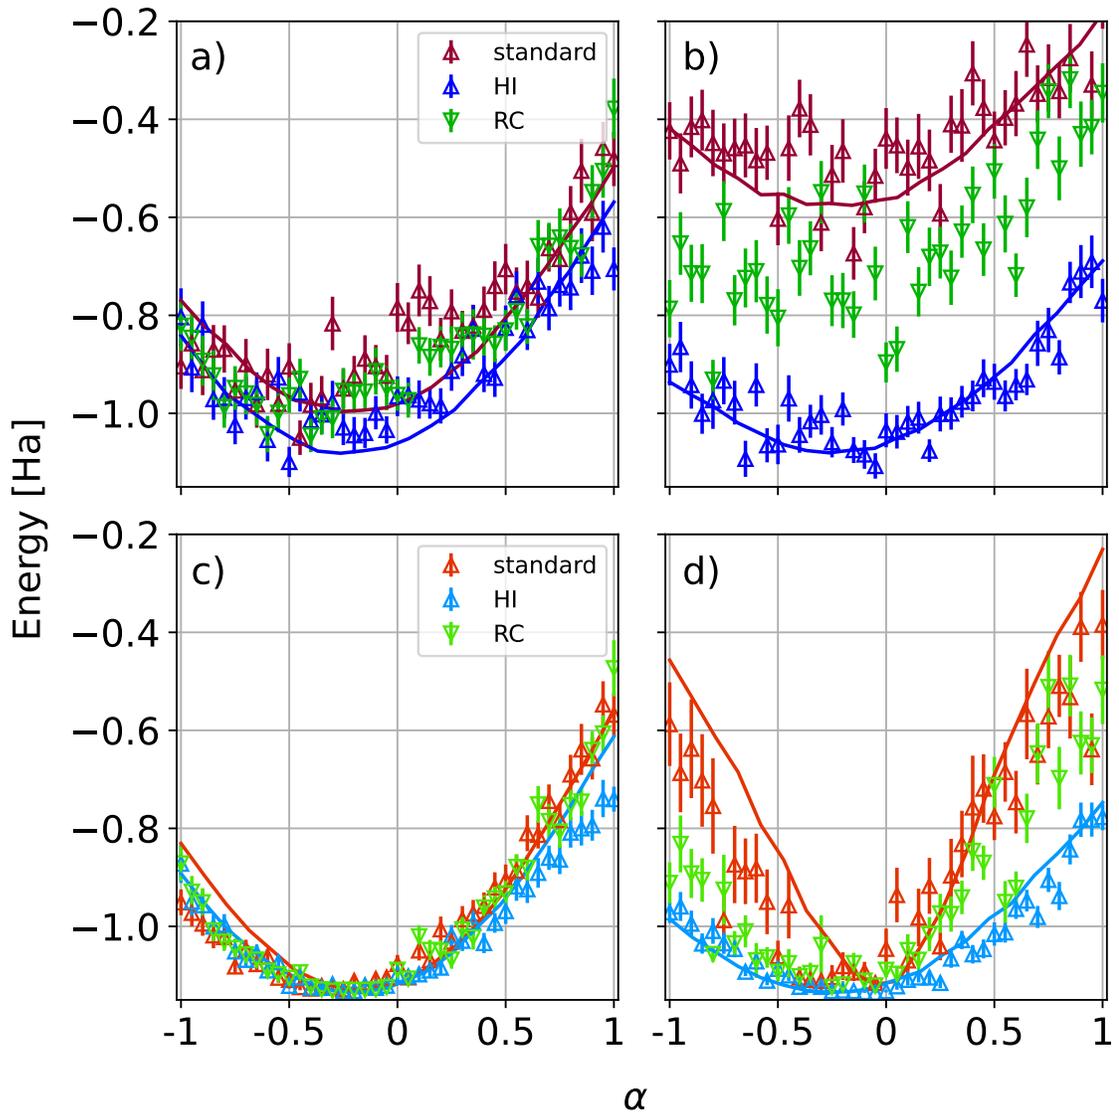


Figure 3.8: **Comparison of HI and RC as mitigation protocols for coherent error** We show raw and simulated (straight lines) energies at ideal calibration (a), energies with 0.5 radian under-rotation injection (b), purified energies at ideal calibration (c), and purified energies 0.5 radian under-rotation injection (d).

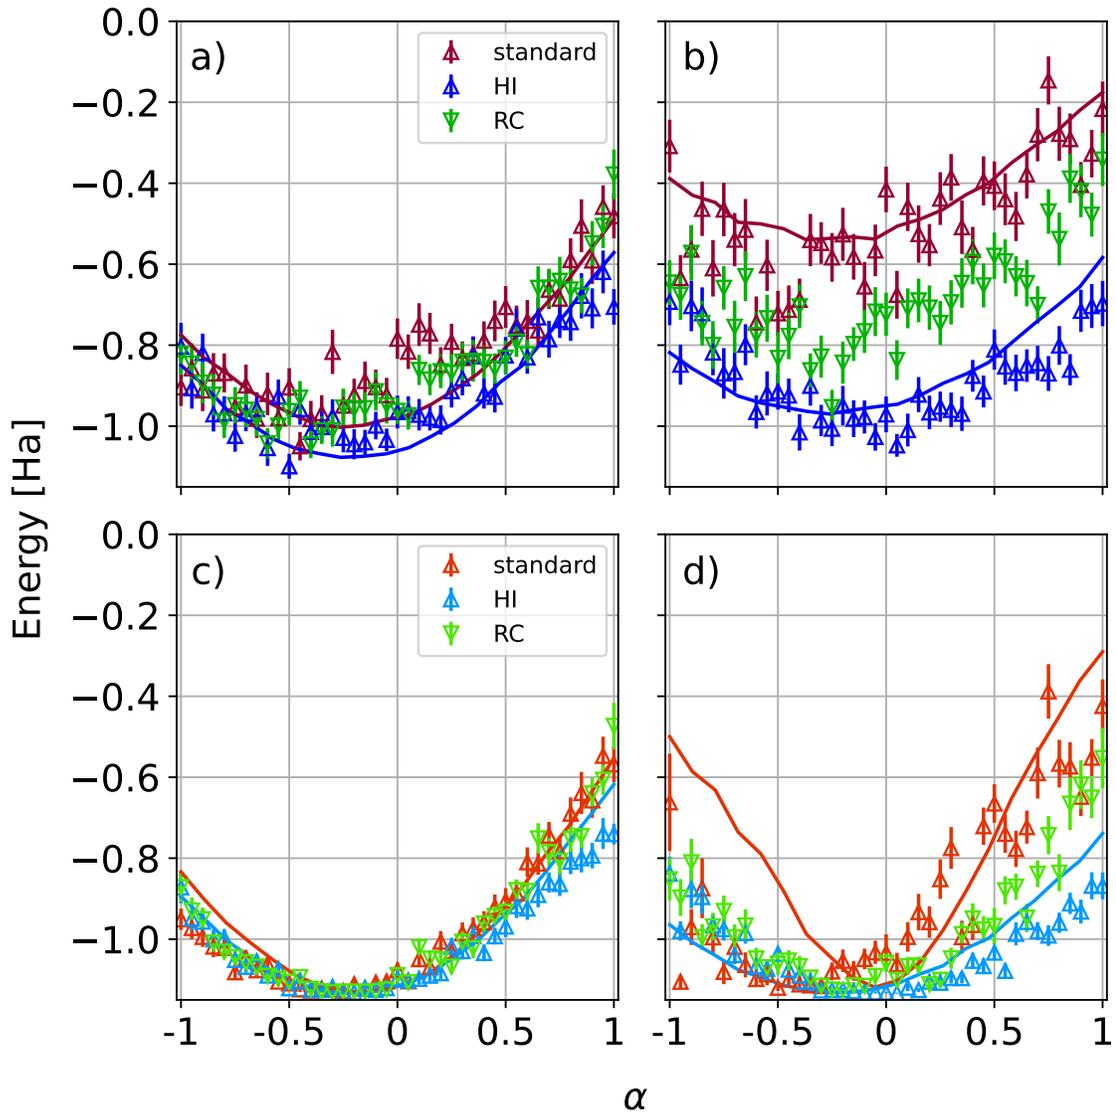


Figure 3.9: **Comparison of HI and RC as mitigation protocols for cooling loop noise** We show raw and simulated energies at ideal calibration with 60 cooling loops (a), energies with only 12 cooling loops (b), purified energies at ideal calibration (c), and purified energies with 12 cooling loops (d).

noise in the experiment further verified the noise model and simulations (and likewise for phase noise), meaning that the noise model adopted here can potentially be used to diagnose potential problems in trapped ion platforms via parameter fitting to a set of data provided by experimentalists. All of the noise mitigation techniques attempted were able to cancel much of the added rotation error as well. Density matrix purification provided a complementary noise mitigation scheme; since it is used in post-processing on classical data provided from measurements, it acts as a “catch-all”, essentially treating coherent and incoherent noise sources indiscriminately. Because of the method’s additional application to stochastic noise and its use in post-processing, it can be effectively used with data provided by either randomly compiled circuits or HI gates. Overall, we find the predictions of our best fit noise model closely align with the experimentally VQE reconstructed loss landscape. Moreover, the parameters of our noise model matches with those estimated in our experimental system. As such, these hidden inverse techniques, in single- and multi-qubit applications, provide a dual purpose, not only error mitigation but also error characterization.

3.7 Hidden inverses in superconducting hardware

In the following, Section 3.7.1, we introduce gate inversion at the pulse level in superconducting devices, followed by gate verification and noise boosting by quantum process tomography and unitary folding experiments. Section 3.7.2 applies the HI to VQE experiments to find the ground state energy of H_2 , and we compare its effectiveness with randomized compilation. In section 3.7.3 we finish with the discussion of our results and numerical simulations and conclusions in Sections 3.7.3 and 3.8, respectively.

3.7.1 Gate inversion on superconducting based architectures

In most of the superconducting architectures, the CX gate is built from an approximated cross-resonance quantum operation $CR = [ZX]_{\pi/2} = \exp(-i\pi\sigma_1^z\sigma_2^x/4)$, introduced in [RD10]. The CR level of approximation defines the quality of CX that can be used on hardware. Thus, it is imperative to find mechanisms to mitigate the level of coherent noise in this gate. The CX gate is locally equivalent to CR by $CX = [ZI]_{-\pi/2}[IX]_{-\pi/2}CR$. A better compilation of CX is used in the IBM quantum device by echoing the CR to mitigate the effects of undesired Hamiltonian terms [She+16b]. Other CX implementation on superconducting devices, that works for any coupled-qubit system with pure transverse coupling, has been proposed using selecting darkening (SD) [De +10; De +12]. The SD method differs from the CR approximation by a single local rotation, therefore, the schemes are locally equivalent [De +12]. Every operation described in the compilation structure represents pulse construction and delivery instructions. We modify these instructions to mimic the inverted gate. In order to illustrate this procedure, we consider a simple example of the π -pulse construction in appendix A. We can generally define the pulse profile throughout the set of amplitudes $\{u_0, \dots, u_N\}$ (waveform) or instructions according to the qiskit library of default functions (parametrized). We extract this information from default gates, and we define a new set of amplitudes $\{-u_N, -u_{N-1}, \dots, -u_0\}$ or new instructions for the parameterized version. In Fig.3.10 (a) and (b), we present the default pulse and its inverse, respectively, for CX_{01} on the IBMQ-Montreal backend. In principle, CX is hermitian, and from the mathematical point of view $CX^{-1} = CX$. Therefore, the new gate should represent the same quantum process. In the experimental setup, we build the CX gate from a pulse structure approximating its quantum operation. We expect that inverting the pulse structure will yield an inverted gate approximation. In other words, for a low

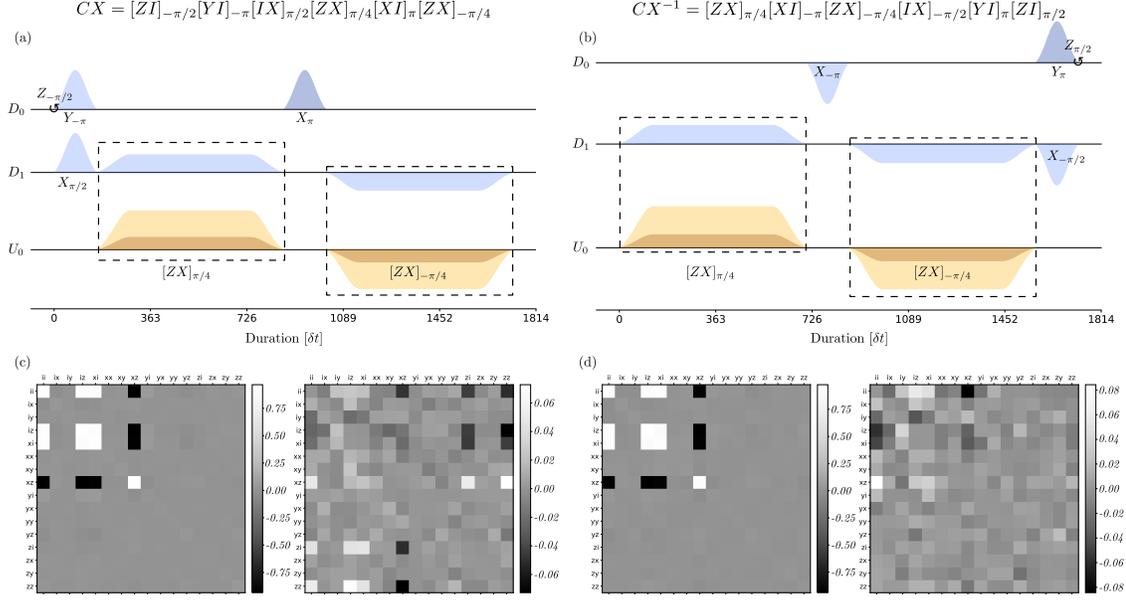


Figure 3.10: **Pulse profile for CX^{-1} and verification:** We implemented a procedure to transform the pulse profile of a given gate, in this case, the CX gate for the qubit pair $[0,1]$ in the IBMQ-Montreal backend, presented in panel (a) into a pulse profile that represents its inverse, presented in panel (b). In addition, we have written the resulting compilation. In panels (c) and (d), we present the process matrix that verifies the operational equivalence between the pulse structure for CX and CX^{-1} .

level of imperfections in the original quantum operation, we expect $CX^{-1} \approx CX$. We verify this assumption by performing quantum process tomography to the gates CX and CX^{-1} and comparing their process matrix (χ -matrix); see Fig. 3.10 (c) and (d). The real part of the χ -matrix for CX^{-1} and CX are similar, indicating that these pulse profiles can be used to represent the same quantum process (up to noise effects). On the other hand, their imaginary parts are different and somewhat statistically relevant ($\max(\chi) > 1/\sqrt{N_S} \sim 10^{-2}$, $N_S = 5000$ the number of shots), which indicates slightly different noise properties. Therefore, the gates represent a similar, albeit not exactly identical, quantum process; it remains to test if the application of the new gate can be used to mitigate noise inside an algorithm on superconducting hardware.

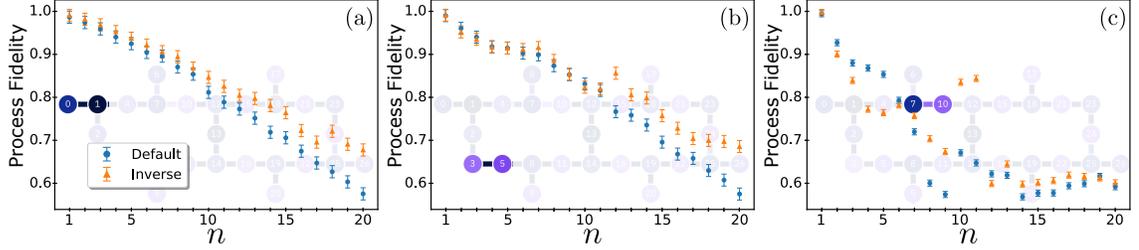


Figure 3.11: **Effective noise scaling on CX gate - IBMQ-Montreal:** This plot presents the unitary folding test in two configurations, using the default gate $\mathcal{U}_0 = CX CX$ (blue circles) and the inverse $\mathcal{U}_1 = CX^{-1}CX$ (orange triangles). Additionally, we consider different backend sectors, showing the noise scaling at different base noise levels. We explored different setups for different values of error rates of local (circle darkness) and nonlocal (line darkness) operations; there, the darker the color of the region in the backend representation, the better is the performance of the quantum operation in terms of fidelity.

Unitary folding test: Next, we consider a slightly different version of the noise scaling introduced [Giu+20] to examine the inverse gate’s dependence on noise injection. In our version, we replace CX with $CX\mathcal{U}^n$, where \mathcal{U} , in theory, is equivalent to the identity operator and n is a positive number. In principle, this replacement should not alter the quantum operation defined by CX and has no logical effect. Now, we choose two versions for \mathcal{U} , a first version $\mathcal{U}_0 = CX CX$ to scale the effective noise on CX , and a second version where $\mathcal{U}_1 = CX^{-1}CX$, which tests the mitigation on CX by applying its inverse. Figure 3.11 shows how the effective noise on CX scales in different sectors of the IBMQ-Montreal backend. We chose three sectors with different one- and two-qubit operational fidelity. Notably, in two of the sectors, noise mitigation is evident by the increase in fidelity as the number of gate repetitions increases.

3.7.2 Application to Variational Quantum algorithms

As a representative algorithm, we consider finding the ground state energy of the H_2 molecule using VQE. This parametrized algorithm updates quantum circuit’s rotational parameters to minimize the Hamiltonian’s expectation value E_θ until it

converges. We consider the 3-parameter approximate unitary coupled-cluster ansatz for four qubits and 2 electrons (referred to as “UCC3” in [McC+19b]), see Fig. 3.12. We systematically invert some of the CX gates in the UCC3 ansatz to cancel the coherent noise present in CX . After every CX applied on qubits (i, j) , we replace the next CX applied on the same pair by its inverse. In Fig. 3.12, we indicate by dashed squares the position of CX^{-1} . For the optimization part of the algorithm, we consider a gradient-free and gradient-based solvers to explore the influence of coherent error mitigation in the learning process under these schemes.

In the experimental implementation, we consider different backends, IBMQ-Bogota, Montreal, and Jakarta, to explore coherent hardware with different noise rates. In addition, we consider low and high-fidelity sectors where the mitigation could be more or less effective. We evaluate HI performance by comparing the VQE learning path with an unmitigated ansatz and a randomly-compiled ansatz with a samples drawn randomly from a set of 20 circuits. We compare both cases to a simulated experiment using the Aer-qiskit simulator. Each experiment used 5000 shots. In Fig. 3.13, we present the learning paths for the above setups.

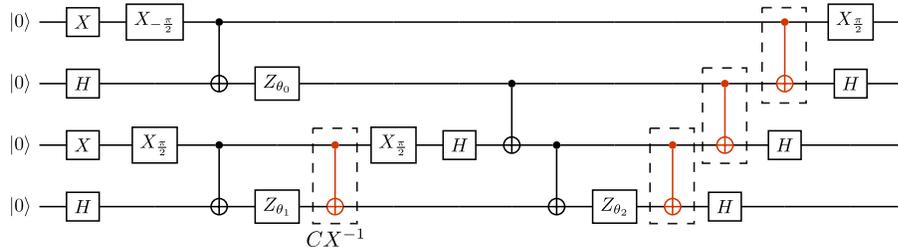


Figure 3.12: **The hidden inverse in the UCC3 ansatz:** The particular arrangement of the CX gates in the UCC3 ansatz allows us to apply their inverse (gates inside the dashed squares) in front of default ones to maximize the coherent noise mitigation.

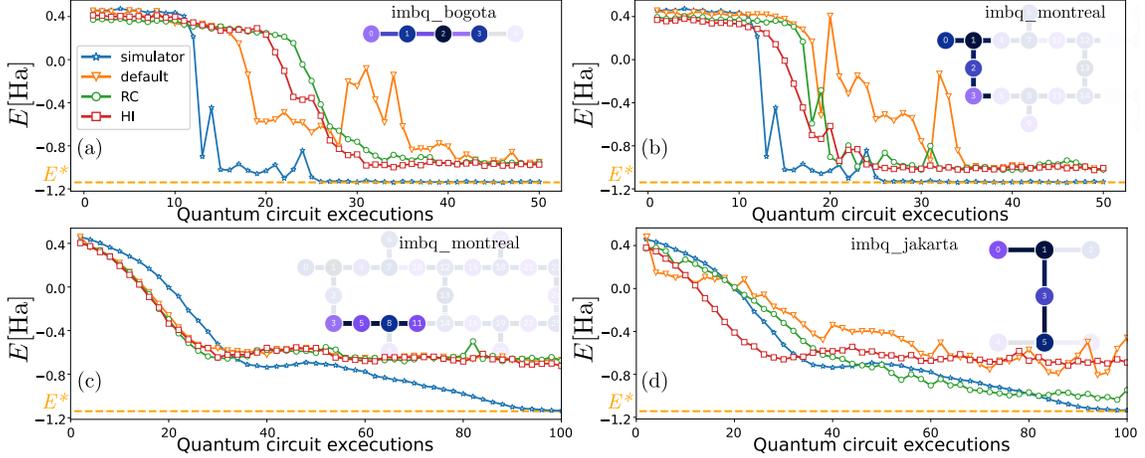


Figure 3.13: **Energy ground state estimation for H_2 using VQE:** In panel (a) using IBMQ-Bogota and (b) IBMQ-Montreal, we present the learning paths with BOBYQA as the optimization algorithm to update the quantum circuit parameters. We present results from randomized compilation (RC), hidden inverse (HI), raw hardware data without QEM (default), and Aer-simulator (simulator) in each panel. In panels (c) and (d) we use the gradient-based solver Adam. These plots show how the energy E approaches the exact value $E^* = -1.1456265$ Ha (orange dashed line). Notably, coherent noise mitigation schemes showed improvement in BOBYQA convergence time, with slightly better performance from the HI circuits, while RC provided higher performance in one case using Adam. In all the experiments, we used 5000 shots. For the randomized compilation, we used 20 random circuit instances.

3.7.3 Discussion

Verification: The real parts of the default and the inverse gates' χ matrices effectively represent the same quantum process to within measurement error. The differences in the imaginary parts, while statistically significant, vary widely from shot to shot, making it difficult to attribute any differences to processes versus background noise. From the experimental point of view, good agreement in the real part of the χ -matrix appears to be enough to provide good results in the unitary folding test.

Unitary folding test: We have verified the operational equivalence between the CX and CX^{-1} with the profile obtained in our procedure. We consider a unitary folding test that consists of the fidelity study of the CX gate when we apply a unitary \mathcal{U}^n for different powers n . As we can expect, the fidelity of the resulting operation

CXU^n decays as n grows. In this version of the unitary folding test, the procedure can amplify incoherent and coherent noise. We consider the coherent noise model

$$CX \rightarrow CX \cdot \mathcal{E} \equiv G_{CX}, \quad (3.4)$$

where the original gate CX is modified by an error map \mathcal{E} producing the operational gate G_{CX} . Under this model, the two folding experiments CXU_0^n and CXU_1 , with $U_0 = G_{CX}G_{CX} = (CX \cdot \mathcal{E})^2$ and $U_1 = G_{CX}G_{CX}^{-1} = \mathbb{I}$. The observed process fidelity shows (see Fig. 3.11(a) and (b)) that we cannot model the operational gate by a product of the original operation (ideal) with an error map. Mainly, the significant fidelity decay in the U_1 folding data suggests an important level of incoherent noise still present in two relevant sectors of IBMQ-Montreal: qubits 0,1 and qubits 3,5, at the time this experiment was performed. Qubit pair 7,10 had the worst reported entangling operation fidelity. As we can see in Fig. 3.11(c), the fidelity fluctuates for U_0 and U_1 folding with no clear evidence of coherence error mitigation for any power n . A possible scenario that explains the inconclusive nature of the unitary folding test on these qubits is that the HI was applied on a non-self-inverse quantum operation. The CX gate between these two qubits, 3 and 7, could be quite different from the ideal CX because our assumed noise model is incorrect. Further, if the default gate is not self-inverse, we cannot expect the HI to mitigate coherent noise because cancellation will not occur.

Application to VQE: The performance of the VQE algorithm, like any other quantum-classical variational algorithm, depends partly on the selection of the classical optimizer and the circuit ansatz. We applied HI to the CX gates used in the UCC3 ansatz, and observed improved convergence performance in the learning process for BOBYQA, a gradient-free solver (see Fig. 3.13 (a) and (b)) in terms of the number of iterations required to reach convergence. When using the Adam optimizer we observe the same performance between HI and RC on Montreal. On

Jakarta, while the HI data converge faster, the RC data converges to a better minimum energy after 100 iterations. The Adam solver is an algorithm for first-order gradient-based optimization with stochastic objective functions. The algorithm is appropriate for non-stationary and very noisy gradients. Therefore, up to certain level of noise, the Adam solver is better adapted to the problem. Thus, the choice of optimal QEM method depends on ansatz, optimizer, and even the underlying qubits. We exploit the sequential application of CX s in the UCC3 ansatz to place CX^{-1} after a CX . The HI requires a CX sequential design, making the method more effective as the sequence gets longer. However, practical considerations also play a role. In general, the application of HI should depend on the magnitude of the rotation in between CX gates; for large rotations, the inverse will amplify rather than cancel coherent noise. In the experiments shown in Fig. 3.13, the nature of batching and queue priority prevented a practical implementation of a conditional rotation check for HI application because it requires a circuit recompilation on a per-shot basis. Therefore the improvements shown for HI in these experiments may not be optimal for all iterations.

To understand the effect of various noise sources and their interaction with the HI during VQE execution, we performed numerical simulations. Specifically, we construct loss landscapes of the VQE example considered above, as loss landscapes can provide key insights into the optimization problem independent of the chosen optimizer. Understanding how noise affects a loss landscape can help us make general predictions about higher level algorithm performance, particularly during optimization. In order to make three-dimensional visualization possible, we fix one of the parameters of the UCC3 ansatz and plot the calculated energy as a function of the two other free parameters. We consider two different noise models in these simulations. First, we consider a time-independent mixed unitary noise channel parameterized by

ϵ (noise strength) and κ (unitarity). This channel is expanded as

$$\varepsilon_G(\rho) = \kappa \cdot \varepsilon_G^c(\rho) + (1 - \kappa) \cdot \varepsilon_G^s(\rho) , \text{ with} \quad (3.5)$$

$$\begin{aligned} \varepsilon_G^c(\rho) &= \exp(-i\epsilon G)\rho \exp(i\epsilon G) \text{ and} \\ \varepsilon_G^s(\rho) &= \cos^2(\epsilon)I\rho I + \sin^2(\epsilon)G\rho G . \end{aligned} \quad (3.6)$$

G is the ideal intended operation, and ε_G is the noisy operation with ε_G^c as the coherent part and ε_G^s as the incoherent part. One nice feature of parameterizing the noise channel this way is that the channel fidelity is independent of κ . This feature allows us to generate comparable coherent and incoherent noise channels by varying κ . Figure 3.14(a) compares the ideal landscape to that of a coherent error channel ($\kappa = 1$ and $\epsilon = 0.02$) and an incoherent error channel ($\kappa = 0$ and $\epsilon = 0.02$). We find that coherent error leads to horizontal shifts in the landscape compared with the ideal case and a slight increase in the minimum energy value.

On the other hand, incoherent error leads to flatter landscapes with a noticeable difference in the minimum energy value. As a result, one can expect a better eigenvalue approximation from coherent errors and a better eigenvector approximation from incoherent errors of similar noise strengths. We also simulated the effect of randomized compiling in the presence of coherent errors, as seen in figure 3.14(b). We find randomized compiling aims to re-center the horizontally shifted landscapes (due to coherent errors) but increases the minimum energy value. In addition, we consider the action of the HI on an induced coherent noise by over-rotation and show how this affects the loss landscape (see Figs. 3.15(a) noisy, and (b) mitigated by HI). We consider a more damaging time-dependent coherent error in the two-qubit entangling operations with $[ZX]_{(1+\epsilon)\theta}$ with $\epsilon \sim \mathcal{N}(0, \varepsilon)$. We calculated a root mean square surface roughness of the loss landscape plot as 0.1065 for the default Ansatz and 0.0611 for the HI-optimized Ansatz. We observe a partial recovery of the loss landscape from the noise influence by applying the hidden-inverse.

This result gives a reasonable inference on the possible reasons for the energy converged values in the on-hardware experiments. Due to incoherent noise, the hidden-inverse and the RC do not improve the converged energy value in the VQE learning path, while they do improve convergence. Applying these methods in combination with another QEM technique, such as readout error mitigation, would potentially yield both faster convergence and a better minimum energy.

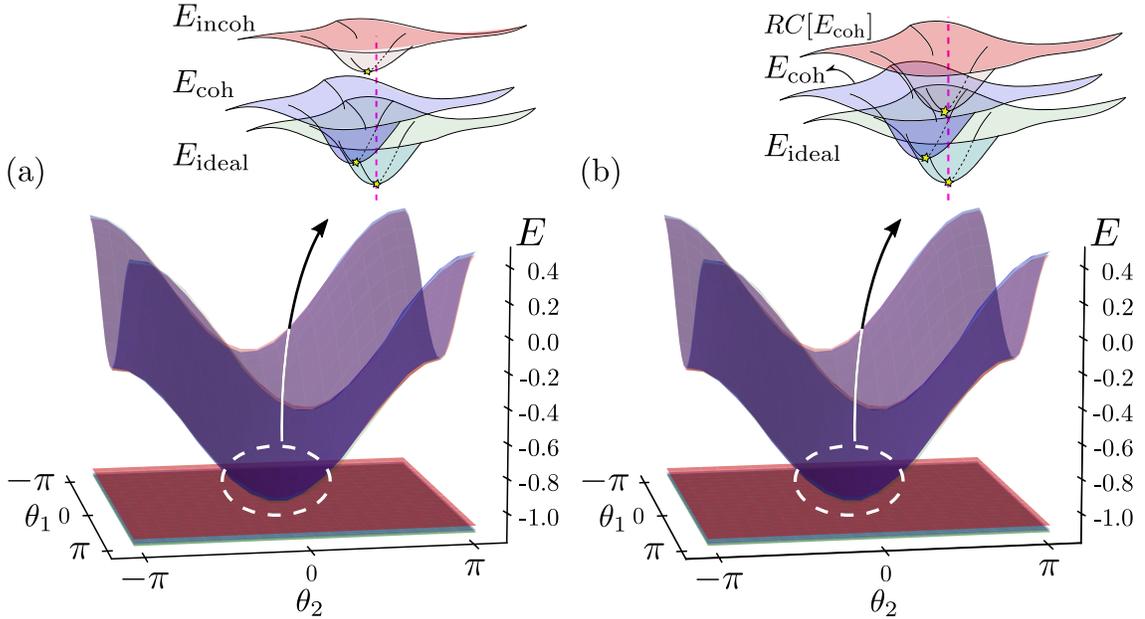


Figure 3.14: Comparing the effect of simulated noise in VQE optimization landscape and the result of applying randomized compiling: We show how a coherent or incoherent error channel modifies the loss landscape of the UCC3 ansatz for the H_2 molecule. (a) Loss landscape for three cases: ideal (green), subjected to coherent noise (blue), and incoherent noise (red). For each case, we draw individual planes representing the minimum of the landscape at the energy levels -1.128 Ha (ideal), -1.1256 Ha (coherent), and -1.098 Ha (incoherent). Inset displays illustration of the landscape behavior (not to scale): coherent error causes horizontal shifts of the landscape and a slight increase in the minimum energy point, while incoherent error results in flattened landscapes with increased minimum energy point. In (b), we present an identical plot as (a), with incoherent error landscapes replaced with the result of applying randomized compiling (red) to the coherent noise channel. We find RC centers the horizontal shifts but raises the minimum energy level (-1.094 Ha).

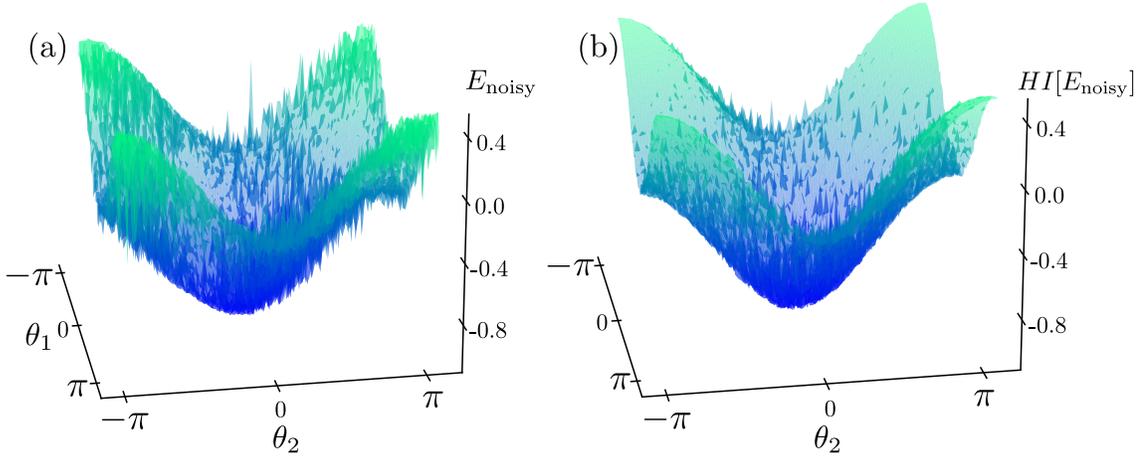


Figure 3.15: **Effect of simulated time-dependent coherent noise in VQE optimization landscape and result of applying HI:** Panel (a) shows the structure of the noisy VQE optimization landscape (E_{noisy}), and panel (b), shows the same landscape with HI applied ($HI[E_{\text{noisy}}]$).

3.8 Conclusions

Here, we have explored methods of coherent noise mitigation and characterization on a trapped ion quantum processor. We note that low-level noise characterization schemes, consisting of both modeling and noise amplification, combined with noise mitigation techniques, can be used to estimate the sources of noise in trapped ion platforms. Low level characterization techniques, beyond solely measuring gate fidelities, help to make a connection between noise at the gate and quantum control levels and higher level performance. Here, we motivated the use of VQE for quantum chemistry as a key high level application whose performance is correlated with the noise sources studied here.

Further, we studied the effects of inverse gates on the performance of a variational quantum algorithm and compared it with randomized compiling. The HI implementation does not require prior knowledge of the noise source besides the coherent noise model and does not require extra quantum resources (gates). We start from the default pulse structure of the gates to build their inverses and apply them

to the latter gate of each pair of CXs in the circuit to enable noise cancellation. The mitigation boosts the performance of BOBYQA, reducing the number of iterations and, thus, the quantum resources required to reach a converged value. We further explored the effectiveness of randomized compiling. We noticed its capability of reducing coherent noise at the cost of additional single qubit gates at each CX location in the circuit. Simulations suggest that incoherent noise is also present in the on-hardware experiments, which matches with prior experience and VQE experiments on superconducting platforms. Incoherent noise could explain the higher ground state found in several cases regardless of the coherent error mitigation method used.

The results show that the HI mitigation mechanism, previously demonstrated on trapped ion platforms, is also effective on superconducting hardware. Combining this QEM method with incoherent noise mitigation schemes should yield improved performance in future VQE applications.

Characterizing large scale quantum systems

In recent years, there has been significant advancements in the development of quantum systems for computing applications. However, as the size of these systems continue to increase, there is a pressing need for practical methods to assess their quality and performance. The traditional approach of fully characterizing an arbitrary quantum process is to estimate its process matrix [NC00; AL04]. However, the number of terms in the process matrix grows exponentially with the system size, making it difficult to accurately estimate the process matrix through protocols such as quantum process tomography or gate set tomography [Nie+20]. This exponential overhead can be circumvented if one is only interested in a quality metric which only requires partial information about the noisy gate. Randomized benchmarking (RB) [EAŻ05; MGE12] experiments can measure the quality of individual gates but have different scalability issues (implementing primitive clifford operations require $\mathcal{O}(N^2/\log N)$ two qubit gates) [EWP+19]. As such RB can only be applied to a few qubits [McK+19] and they cannot detect system errors that appear only when one scales up the system size. System errors such as crosstalk can have a significant impact on the performance of large-scale systems. To address these limitations, techniques such as quantum volume [Cro+19] and cross-entropy benchmarking [Boi+18] have been proposed to certify the performance of large-scale quantum systems. While these protocols have been successful in the short term, they are not truly scalable as they require access to a ground truth that is not computationally feasible to simulate

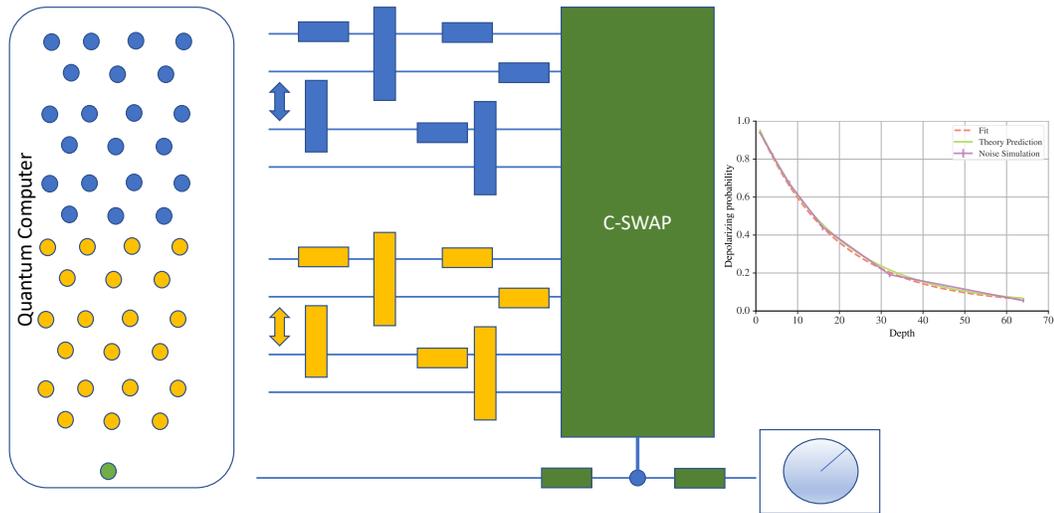


Figure 4.1: **Estimating error per circuit layer.** We breakdown the number of qubits to characterize into two parts, run identical random circuits on both and finally perform a swap test.

as the system size scales.

In this chapter, we present a scalable error characterization protocol that aims to certify the performance of large-scale quantum computers without the need for classical simulation. To this end, we propose a new protocol Error Per Circuit Layer (EPCL), to estimate the error in a layer of quantum operation independent of the system size, and take into account various local and global error sources while being SPAM-resistant. The metric produced by this protocol represents the depolarizing probability of a global depolarizing channel acting on the entire qubit register after a single layer of circuit operations. Additionally, our protocol can be used to detect systematic noise, spatial homogeneity of noise, and the weight of error processes.

4.1 Estimating average fidelity of cycles

We now introduce the protocol to estimate device noise of a quantum processor (or between multiple processors). It can be seen as combining purity estimation [WE16] within a randomized benchmarking framework. We assume the device has an incoherent independent noise channel. Our protocol can be used to estimate fidelity of cycles of quantum operations sampled from universal gate set without requiring any classical simulation.

- Identify two ensembles of n qubit quantum processor to characterize and a single ancilla qubit for a total of $2n + 1$ qubits.
- Select a set of depths (cycles) $d = (1, \dots, d_{max})$.
- For each depth d_i , generate M random sequences on n qubits, drawing gates from an universal gate set.
- Apply the same sequence on the two n -qubit ensembles, and perform a swap test on the entire register with the measurement qubit as the control qubit in the swap test. Repeat this procedure for different sequences and over all depths and store the state overlap $2 \times (P_{m=0,d_i} - \frac{1}{2})$ where $P_{m=0,d_i}$ is probability of measuring the ancilla qubit in zero state for depth d_i . In the absence of noise this quantity is always one.
- Fit $\mathbb{E} [2 \times (P_{m=0,d_i} - \frac{1}{2})] = \frac{1}{2^n} - p^{2d_i} \frac{(1-A \times 2^n)}{2^n}$ where p is the depolarizing probability per layer.
- Estimate p .

4.2 Algorithmic modification in the presence of coherent error

In the presence of spatially homogeneous coherent errors in the device, we can modify the protocol to incorporate these errors in the noise estimate. The modification is made in Step 4, where instead of applying the same circuit to two n-qubit ensembles, we generate multiple randomly compiled versions of the circuit [WE16]. These randomly compiled versions are logically equivalent in the absence of noise. By applying two different randomly compiled versions, we can effectively twirl the noise. From there, we can proceed with the rest of the procedure outlined in Section 4.2.

4.3 Numerical simulations and Experimental implementations

4.3.1 Noise model for numerical simulations

In order to numerically simulate and verify our protocol, we use different noise models such as depolarizing noise, realistic mixed coherent incoherent noise, interaction errors like cross talk, and error from circuit mapping overhead of controlled swap operations. In this section, we describe these models and present the outcomes of our protocol.

Depolarizing noise: Our simplest model is a depolarizing noise channel on single and two qubit gates. In our simulation we use the same depolarizing probability, $p_d=0.01$, for single qubit and two qubit gates. This is because our protocol does not distinguish between single qubit and two qubit errors but we are interested in the total error of a layer. This depolarizing probability corresponds to 99.5% single qubit gate fidelity and 99.25% two qubit gates. We use a 11-qubit processor, to characterise the noise after a layer of quantum operations on 5-qubits. To do so, we divide the processor into two sets of 5-qubits and an ancilla qubit, and apply our protocol. The data from the simulation, as well as the corresponding fit is plotted in Fig. 4.2. The estimated 5 qubit depolarizing channel has a depolarizing probability (\tilde{p}_d^5) of

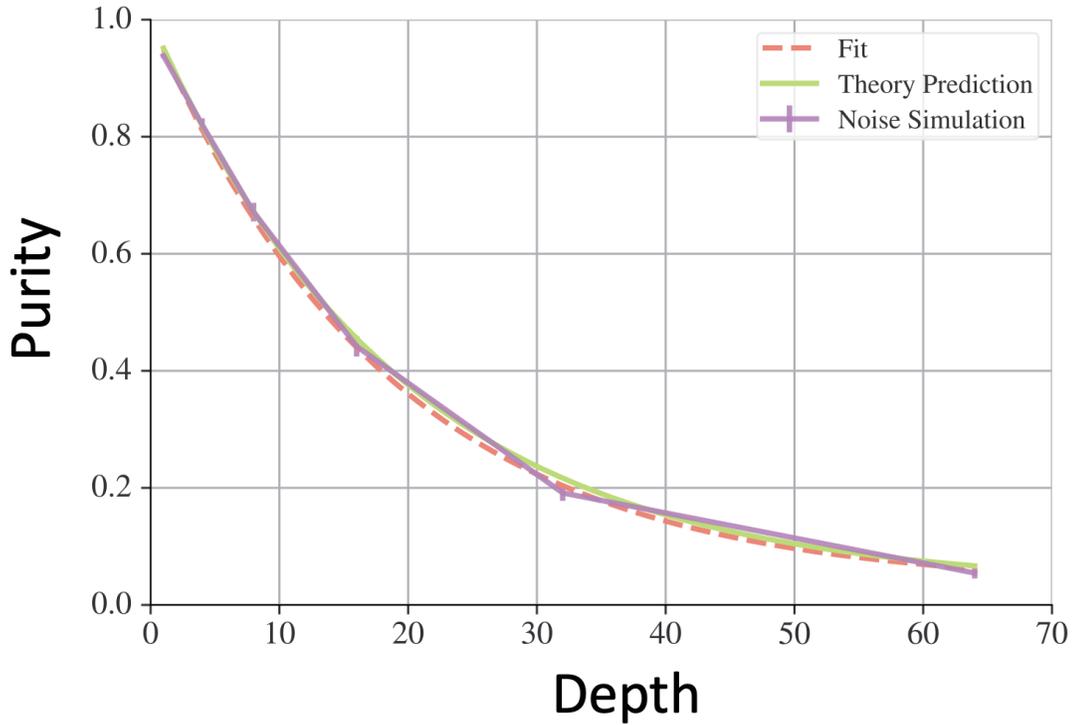


Figure 4.2: **Purity plotted against the number of layers for a noise model with single qubit and two qubit depolarizing noise.**

0.026 which results in a per layer fidelity of 97.48%. If we assume the noise channel is a single qubit depolarizing noise on every qubit per layer, then the equivalent per qubit fidelity is 99.49%.

Mixed unitary channel: A more physically realistic model involves a mixed coherent incoherent and incoherent channel parameterized by two parameters ϵ and κ . ϵ dictates the strength of the channel and κ reflects the coherence of the noise. We can model the behavior of the noise process as ideal gate followed by this noise channel [Wal+15; She+16a; Fen+16; Deb+18]

$$\varepsilon_G(\rho) = \kappa \cdot \varepsilon_G^c(\rho) + (1 - \kappa) \cdot \varepsilon_G^s(\rho) , \text{ with} \quad (4.1)$$

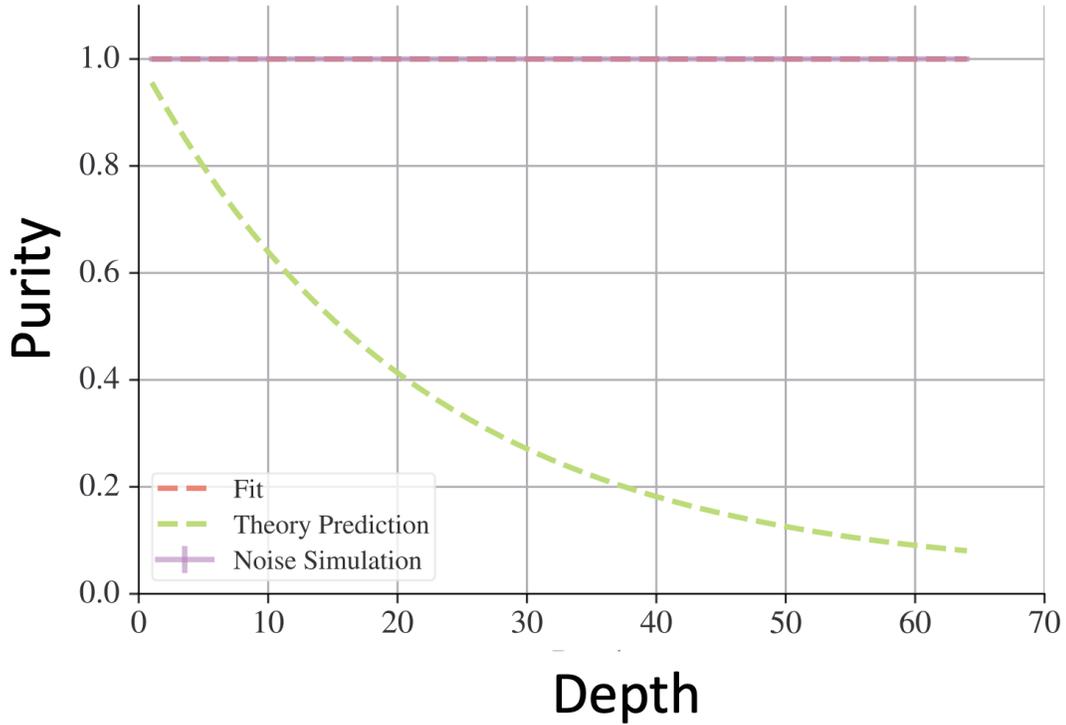


Figure 4.3: **Purity plotted against the number of layers for a noise model with purely coherent error.** We find the estimated probabilities do not match with the theoretical prediction as the noise is identical and coherent in both circuits.

$$\begin{aligned}
 \varepsilon_G^c(\rho) &= \exp(-i\epsilon G)\rho \exp(i\epsilon G) \text{ and} \\
 \varepsilon_G^s(\rho) &= \cos^2(\epsilon)I\rho I + \sin^2(\epsilon)G\rho G .
 \end{aligned}
 \tag{4.2}$$

In our numerical simulations, we fix $\epsilon = 0.05$ and $\kappa = 1$ which results in a purely coherent channel and demonstrate the need for modifying our algorithm as in Sec 4.2 to account for coherent errors. We plot the result from the simulation in Fig 4.3, which shows that our original protocol fails. This is because if the noise is identical and coherent on both of the circuits, the resulting density matrix is pure hence the purity does not decay as we increase circuit depth. But instead if we use the modified protocol where we apply randomly compiled circuits on the two sets of qubits, we get Fig 4.4. For the simulation here, we pick 100 different randomly compiled circuits. The exponential decay in Fig 4.4 shows that the noise channel has been twirled. The

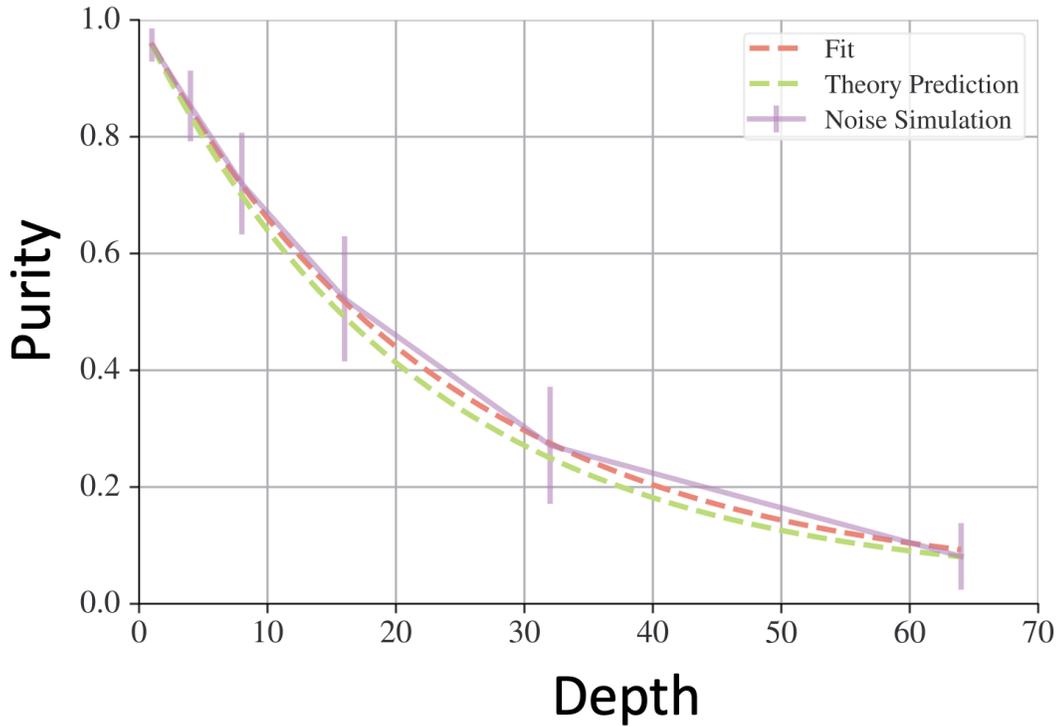


Figure 4.4: **Purity plotted against the number of layers for a noise model with purely coherent error but with modified protocol that twirls the error independently on the two circuits.** We find the estimated probabilities now do match with the theoretical prediction.

theoretical fidelity of a single layer across five qubits in with the noise parameters is 97.59% with depolarizing probability of 0.025 and the estimated single layer fidelity from our protocol is 97.48% and depolarizing probability is 0.026. As seen in Fig 4.4, there is a good match between the theoretical prediction and noisy simulation.

System errors: Crosstalk

System errors in quantum computers are errors that arise due to the complex interactions between the many quantum bits (qubits) that make up a quantum computer. These errors can be difficult to predict and can have a significant impact on the accuracy of quantum computations. One example of such an interaction error is crosstalk. Crosstalk in quantum computers refers to the phenomenon where the

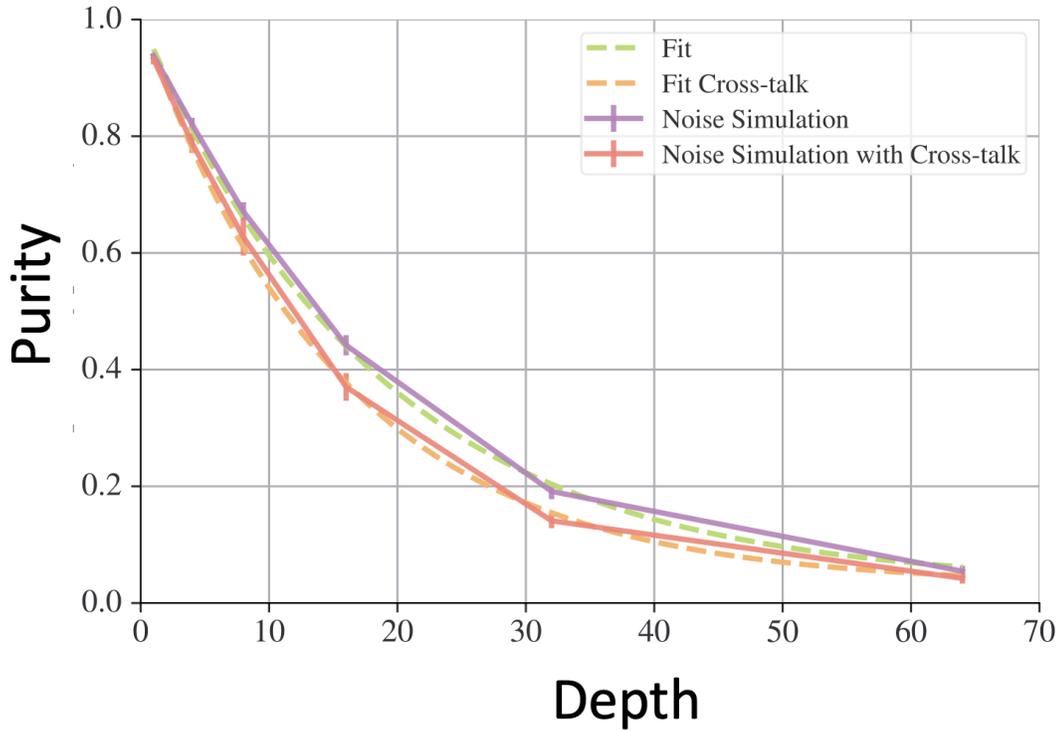


Figure 4.5: **Comparison of purity plotted against the number of layers for a noise model with a cross-talk and one without.** We find the estimated probabilities are sensitive to cross-talk error.

operation of one qubit can affect the state of another qubit. This can happen when qubits are placed too close to each other, or have similar addressing frequency allowing them to interact and interfere with each other. We simulated two noise models one with and the other one without crosstalk and demonstrate how our protocol is sensitive to this interaction error as expected. Our noise model only includes nearest neighbor crosstalk for simplicity and crosstalk is modeled as depolarizing noise in neighboring ions. In reality, just like gate errors, crosstalk will also have a coherent part which can be twirled. Fig 4.5 shows the simulated results and the fit with and without crosstalk error. We find the fidelity per layer without crosstalk is 97.43% and fidelity with crosstalk is 96.95%.

Swap gate errors: One apparent challenge of this protocol is the high cost of

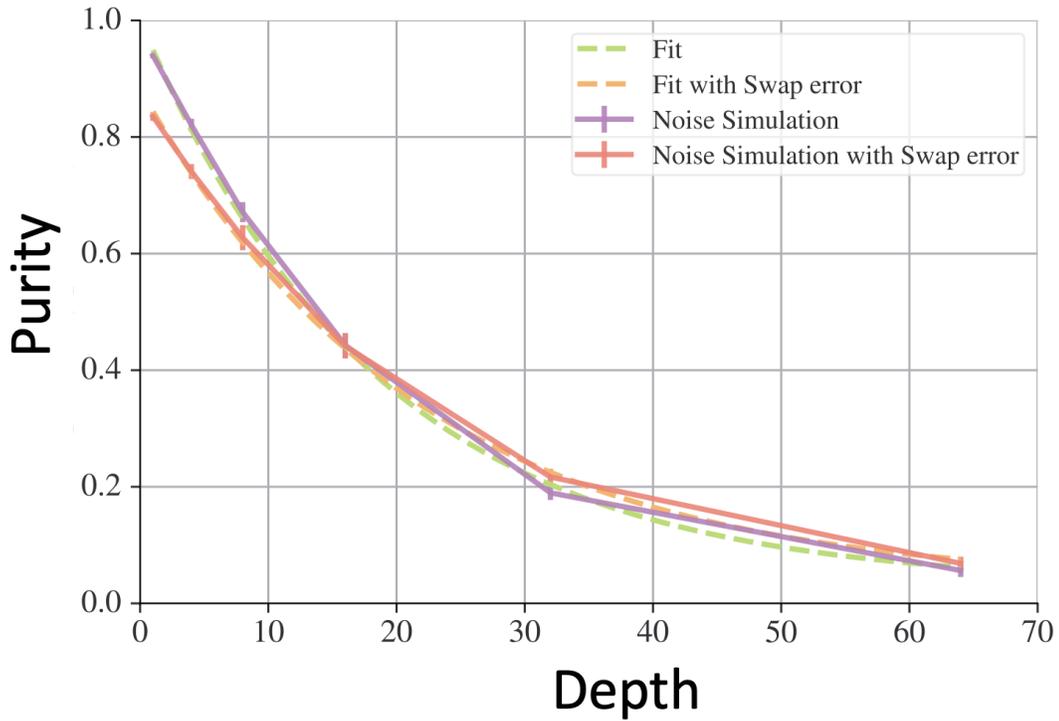


Figure 4.6: **Effect of noisy swap test on the performance of the protocol.** We find that noisy swap gates do reduce resolution but they can be seen as SPAM error.

implementing a controlled swap gate. But this not a big problem as the error in the swap gate can be absorbed in the state preparation and measurement error. In Fig 4.6, we plot the decaying plots with and without added noise in the swap gate. We find the added swap error reduces the resolution of the plots as is common in randomized benchmarking protocols but the estimated fidelity of the operations are mostly unaffected. This does however put a limit on how big of a system can be characterized for a given gate error rate. One way of getting around this challenge is to replace the swap test with a pairwise Bell basis measurement and classical processing on the entire register sometimes also known as a destructive swap test [GC13].

4.3.2 Testing for spatial in-homogeneity

In recent years, there has been significant progress in the development of quantum computers with more than one hundred qubits. One approach to scaling up to even larger systems is based on a modular architecture, where a group of smaller quantum computers are interconnected through quantum communication. While the individual blocks can be characterized independently, it is also important to understand the level of homogeneity of noise across the different blocks. The protocol described in this work can be used to assess the degree of inhomogeneity between different modules in a distributed quantum computing setup.

4.4 Derivation of the fit model

4.4.1 Estimating purity from Swap test

Swap test [Buh+01] is a procedure to compute the overlap between two quantum states. One way to implement the swap test is to apply a controlled swap operation as depicted on Fig 4.7. Swap test take two inputs as $|\psi\rangle$ and $|\phi\rangle$. Then the probability of measuring the ancilla qubit to be in 0 is $\frac{1}{2} + \frac{1}{2}|\langle\psi|\phi\rangle|^2$. Instead of pure states, if we have two mixed states (ρ, σ) swap test gives us the overlap $\frac{1}{2} + \frac{1}{2}\text{Tr}(\rho\sigma)$. If the two input states are the same, we finally get the purity of the input state $\frac{1}{2} + \frac{1}{2}\text{Tr}(\rho^2)$.

4.4.2 Error scrambling under random circuits

In this protocol, we apply layers of random quantum circuit on an initial state. The effect of applying noisy random quantum circuit is modeled as an unitary transformation followed by an n qubit depolarizing channel with depolarizing probability p. We can write the resulting state as

$$|\psi\rangle \rightarrow \rho_U = p|\psi_U\rangle\langle\psi_U| + (1-p)I/2^n$$

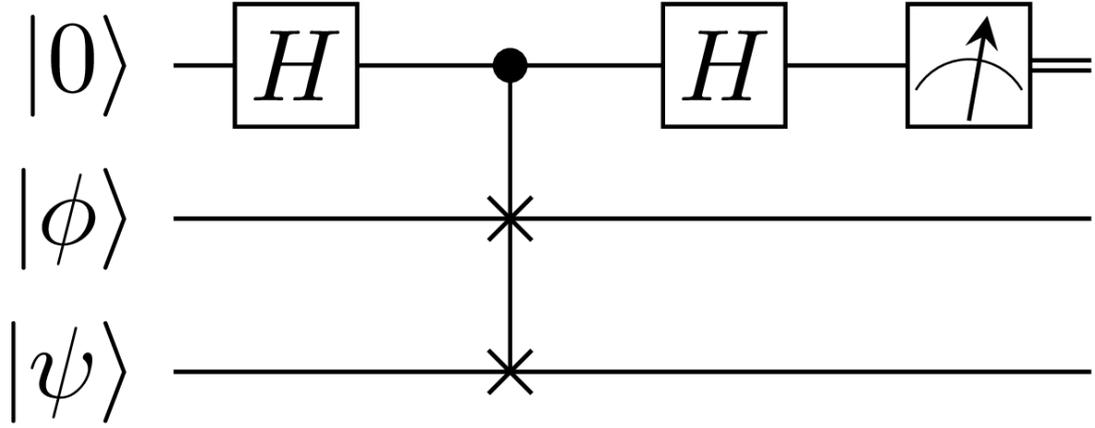


Figure 4.7: **Swap test circuit**

where $|\psi_U\rangle$ is the ideal output state which is a result of applying noise-free random quantum circuit on some (pure) initial state. If we apply d layers of the random quantum circuit, the resulting state can be written as

$$|\psi\rangle \rightarrow \rho_U = p^d |\psi_U\rangle \langle\psi_U| + (1 - p^d)I/2^n$$

. Now, if we have two copies of this state, and we perform a swap test on it, the probability of measuring the ancilla in the ground state is $\frac{1}{2} + \frac{1}{2}Tr(\rho_U^2)$

4.4.3 Estimating depolarizing probability and fidelity

Let us calculate $Tr(\rho_U^2)$ for our state ρ_U . For simplicity, we will write $|\psi_U\rangle$ as $\tilde{\rho}_U$:

$$\begin{aligned} Tr(\rho_U^2) &= Tr((p^d \tilde{\rho}_U + (1 - p^d)I/2^n)(p^d \tilde{\rho}_U + (1 - p^d)I/2^n)) \\ &= Tr(p^{2d} \tilde{\rho}_U^2 + 2p^d(1 - p^d)I/4^n + (1 - p^d)^2 I/4^n = \\ & \qquad \qquad \qquad 1/2^n - p^{2d}(1 - A2^n)/2^n \end{aligned} \tag{4.3}$$

where $A = Tr(\tilde{\rho}^2)$ and we have used $Tr(I) = 2^n$. In the absence of SPAM errors, $Tr(\tilde{\rho}^2)$ is 1 as $\tilde{\rho}_U$ is a pure state. We can use this fit to extract of the depolarizing probability. Once we have calculated the n-qubit depolarizing probability, we can calculate the fidelity of the n-qubit layer as $\mathcal{F}_n = (1 - p) + p/2^n$

Conclusion

The research presented in this thesis focused on the development of techniques to evaluate and control noise in quantum systems. These methods not only have relevance to near-term quantum systems but also for future fault-tolerant systems. One of the primary challenges faced in quantum systems is time-dependent errors, which can severely affect performance. In this work, we propose closed-loop feedback based quantum control protocols that can correct for these errors, leading to longer and more accurate computations. This newly introduced characterization protocol has recently been demonstrated to be effective in practice. As reported in [Sin+22], recent experimental results show that the protocol has improved the correction of correlated phase errors in an array of rubidium data qubits.

Moreover, the use of quantum error mitigation protocols can pave the way for the near-term application of quantum computers before the development of fault-tolerant systems. Our contribution in this area includes the introduction of the concept of hidden inverses, which proves effective in reducing errors in entangling gates without incurring additional quantum resource costs. Our experiments on various platforms demonstrate the improvement in variational algorithms from using hidden inverses.

Finally, as quantum systems become larger, it becomes imperative to develop effective characterization techniques. Existing benchmarking methods for large quantum systems rely on access to ideal classical simulations, which are not scalable, limiting the size of the quantum computer that can be evaluated. To overcome this

limitation, we introduce a new characterization protocol that provides the error per circuit layer of operations and does not rely on classical simulations.

Bibliography

- [AB97] D. Aharonov and M. Ben-Or. “Fault-tolerant Quantum Computation with Constant Error”. In: *Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing*. STOC '97. El Paso, Texas, USA: ACM, 1997, pp. 176–188. ISBN: 0-89791-888-6. DOI: 10.1145/258533.258579.
- [AC07] Panos Aliferis and Andrew W. Cross. “Subsystem Fault Tolerance with the Bacon–Shor Code”. In: *Phys. Rev. Lett.* 98 (22 May 2007), p. 220502. DOI: 10.1103/PhysRevLett.98.220502.
- [ADR82] Alain Aspect, Jean Dalibard, and Gérard Roger. “Experimental Test of Bell’s Inequalities Using Time-Varying Analyzers”. In: *Phys. Rev. Lett.* 49 (25 Dec. 1982), pp. 1804–1807. DOI: 10.1103/PhysRevLett.49.1804. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.49.1804>.
- [AL04] G. M. D. Ariano and P. Lo Presti. “Quantum State Estimation”. In: *Lecture Notes in Physics*. Ed. by M. Paris and J. Rehcek. Vol. 649. Berlin: Springer-Verlag, 2004, p. 297.
- [Asa+19] Serwan Asaad, Vincent Mourik, Benjamin Joecker, et al. “Coherent electrical control of a single high-spin nucleus in silicon”. In: *Preprint at https://arxiv.org/abs/1906.01086* (2019).
- [AT17] M. P. Allen and D.J. Tildesley. “Computer Simulations of Liquids”. In: *Oxford University Press* (2017).
- [AW59] B. J. Alder and T. E. Wainwright. “Studies in Molecular Dynamics. I. General Method”. In: *J. Chem. Phys.* 31 (1959), p. 459. DOI: 10.1063/1.1730376.
- [Bal+16] C. J. Ballance, T. P. Harty, N. M. Linke, et al. “High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits”. In: *Phys. Rev. Lett.* 117 (6 Aug. 2016), p. 060504. DOI: 10.1103/PhysRevLett.117.060504. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.060504>.

- [Ban+13] Masamitsu Bando, Tsubasa Ichikawa, Yasushi Kondo, et al. “Concatenated Composite Pulses Compensating Simultaneous Systematic Errors”. In: *J. Phys. Soc. Jpn.* 82.1 (2013), p. 014004. DOI: 10.7566/JPSJ.82.014004.
- [BB19] Natalie C. Brown and Kenneth R. Brown. “Leakage mitigation for quantum error correction using a mixed qubit scheme”. In: *Phys. Rev. A* 100 (3 Sept. 2019), p. 032325. DOI: 10.1103/PhysRevA.100.032325.
- [BB48] J. Bardeen and W. H. Brattain. “The Transistor, A Semi-Conductor Triode”. In: *Phys. Rev.* 74 (2 July 1948), pp. 230–231. DOI: 10.1103/PhysRev.74.230. URL: <https://link.aps.org/doi/10.1103/PhysRev.74.230>.
- [BB84] Charles H. Bennett and Gilles Brassard. “Quantum cryptography: Public key distribution and coin tossing”. In: (1984), pp. 175–179.
- [BC94] Samuel L. Braunstein and Carlton M. Caves. “Statistical distance and the geometry of quantum states”. In: *Phys. Rev. Lett.* 72 (22 May 1994), pp. 3439–3443. DOI: 10.1103/PhysRevLett.72.3439.
- [Bec+03] J. Becker, G. Grun, R. Seemann, et al. “Complex dewetting scenarios captured by thin-film models”. In: *Nat. Mater.* 2.1 (Jan. 2003), pp. 59–63.
- [Ben+82] Roberto Benzi, Giorgio Parisi, Alfonso Sutera, et al. “Stochastic resonance in climatic change”. In: *Tellus* 34.1 (1982), pp. 10–16.
- [Ber+22] Ewout van den Berg, Zlatko K Mineev, Abhinav Kandala, et al. “Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors”. In: *arXiv preprint arXiv:2201.09866* (2022).
- [BHC04] Kenneth R. Brown, Aram W. Harrow, and Isaac L. Chuang. “Arbitrarily accurate composite pulse sequences”. In: *Phys. Rev. A* 70 (5 Nov. 2004). K. R. Brown, A. W. Harrow, and I. L. Chuang. *Phys. Rev. A* 72, 039005 (2005)., p. 052318. DOI: 10.1103/PhysRevA.70.052318.
- [Blu+13] Robin Blume-Kohout, John King Gamble, Erik Nielsen, et al. “Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit”. In: *arXiv preprint arXiv:1310.4492* (2013). URL: <https://arxiv.org/abs/1310.4492> (visited on 07/15/2017).
- [Boi+18] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, et al. “Characterizing quantum supremacy in near-term devices”. In: *Nature Physics* 14 (6 Apr. 2018), p. 595. DOI: 10.1038/s41567-018-0124-x. URL: <https://doi.org/10.1038/s41567-018-0124-x>.
- [Bol+96] J. J. Bollinger, Wayne M. Itano, D. J. Wineland, et al. “Optimal frequency measurements with maximally correlated states”. In: *Phys. Rev. A* 54 (6 Dec. 1996), R4649–R4652. DOI: 10.1103/PhysRevA.54.R4649.

- [Bon+18] X Bonet-Monroig, R Sagastizabal, M Singh, et al. “Low-cost error mitigation by symmetry verification”. In: *Phys. Rev. A* 98.6 (2018), p. 062339.
- [Bor26] Max Born. “Zur Quantenmechanik der Stoßvorgänge”. In: *Zeitschrift für Physik* 37.12 (1926), pp. 863–867.
- [Buh+01] Harry Buhrman, Richard Cleve, John Watrous, et al. “Quantum Fingerprinting”. In: *Phys. Rev. Lett.* 87 (16 Sept. 2001), p. 167902. DOI: 10.1103/PhysRevLett.87.167902. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.87.167902>.
- [BV97] Ethan Bernstein and Umesh Vazirani. “Quantum Complexity Theory”. In: *SIAM Journal on Computing* 26.5 (1997), pp. 1411–1473.
- [Byl+11] Jonas Bylander, Simon Gustavsson, Fei Yan, et al. “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit”. In: *Nat. Phys.* 7 (2011), p. 565.
- [Che+19] Yanzhu Chen, Maziar Farahzad, Shinjae Yoo, et al. “Detector tomography on IBM quantum computers and mitigation of an imperfect measurement”. In: *Physical Review A* 100.5 (2019), p. 052315.
- [Cla+21] Susan M. Clark, Daniel Lobser, Melissa C. Revelle, et al. “Engineering the Quantum Scientific Computing Open User Testbed”. In: *IEEE Transactions on Quantum Engineering* 2 (2021), pp. 1–32. DOI: 10.1109/TQE.2021.3096480.
- [Cla+69] John F. Clauser, Michael A. Horne, Abner Shimony, et al. “Proposed Experiment to Test Local Hidden-Variable Theories”. In: *Phys. Rev. Lett.* 23 (15 Oct. 1969), pp. 880–884. DOI: 10.1103/PhysRevLett.23.880. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.23.880>.
- [Col+19] The Event Horizon Telescope Collaboration, Kazunori Akiyama, Anton Alberdi, et al. “First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole”. In: *The Astrophysical Journal Letters* 875.1 (Apr. 2019), p. L1. DOI: 10.3847/2041-8213/ab0ec7. URL: <https://dx.doi.org/10.3847/2041-8213/ab0ec7>.
- [Cro+19] Andrew W. Cross, Lev S. Bishop, Sarah Sheldon, et al. “Validating quantum computers using randomized model circuits”. In: *Phys. Rev. A* 100 (3 Sept. 2019), p. 032328. DOI: 10.1103/PhysRevA.100.032328. URL: <https://link.aps.org/doi/10.1103/PhysRevA.100.032328>.
- [CRR16] I. Cohen, A. Rotem, and A. Retzker. “Refocusing two-qubit-gate noise for trapped ions by composite pulses”. In: *Phys. Rev. A* 93 (3 Mar. 2016), p. 032340. DOI: 10.1103/PhysRevA.93.032340.

- [CZ95] J. I. Cirac and P. Zoller. “Quantum Computations with Cold Trapped Ions”. In: *Phys. Rev. Lett.* 74 (20 May 1995), pp. 4091–4094. DOI: 10.1103/PhysRevLett.74.4091. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.74.4091>.
- [Cza+20] Piotr Czarnik, Andrew Arrasmith, Patrick J Coles, et al. “Error mitigation with Clifford quantum-circuit data”. In: *arXiv preprint arXiv:2005.10189* (2020).
- [Cza+21] Piotr Czarnik, Andrew Arrasmith, Patrick J. Coles, et al. “Error mitigation with Clifford quantum-circuit data”. In: *Quantum* 5 (Nov. 2021), p. 592. ISSN: 2521-327X. DOI: 10.22331/q-2021-11-26-592. URL: <https://doi.org/10.22331/q-2021-11-26-592>.
- [DCS17] Rafał Demkowicz-Dobrzański, Jan Czapkowski, and Pavel Sekatski. “Adaptive Quantum Metrology under General Markovian Noise”. In: *Phys. Rev. X* 7 (4 Oct. 2017), p. 041009. DOI: 10.1103/PhysRevX.7.041009.
- [De +10] PC De Groot, J Lisenfeld, RN Schouten, et al. “Selective darkening of degenerate transitions demonstrated with two superconducting quantum bits”. In: *Nature Physics* 6.10 (2010), pp. 763–766.
- [De +12] PC De Groot, S Ashhab, A Lupaşcu, et al. “Selective darkening of degenerate transitions for implementing quantum controlled-NOT gates”. In: *New Journal of Physics* 14.7 (2012), p. 073038.
- [Deb+16] S. Debnath, N. M. Linke, C. Figgatt, et al. “Demonstration of a small programmable quantum computer with atomic qubits”. In: *Nature* 536.7614 (Aug. 2016), pp. 63–66. ISSN: 1476-4687. DOI: 10.1038/nature18648. URL: <https://doi.org/10.1038/nature18648>.
- [Deb+18] Dripto M. Debroy, Muyuan Li, Michael Newman, et al. “Stabilizer Slicing: Coherent Error Cancellations in Low-Density Parity-Check Stabilizer Codes”. In: *Phys. Rev. Lett.* 121 (25 Dec. 2018), p. 250502. DOI: 10.1103/PhysRevLett.121.250502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.250502>.
- [Des+04] L. Deslauriers, P. C. Haljan, P. J. Lee, et al. “Zero-point cooling and low heating of trapped $^{111}\text{Cd}^+$ ions”. In: *Phys. Rev. A* 70 (4 Oct. 2004), p. 043408. DOI: 10.1103/PhysRevA.70.043408. URL: <https://link.aps.org/doi/10.1103/PhysRevA.70.043408>.
- [Deu92] David Deutsch. “Rapid solution of problems by quantum computation”. In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 439.1907 (1992), pp. 553–558.
- [DG28] C. J. Davisson and L. H. Germer. “Reflection of Electrons by a Crystal of Nickel”. In: *Proceedings of the National Academy of Sciences of the United States of America* 14.4 (1928), pp. 317–322.

- [Dir28] P. A. M. Dirac. “The Quantum Theory of the Electron”. In: *Proceedings of the Royal Society A* 117.778 (1928), pp. 610–624.
- [Dir30] P. A. M. Dirac. “Note on Exchange Phenomena in the Thomas Atom”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 26.3 (1930), pp. 376–385. DOI: 10.1017/S0305004100016108.
- [EAŻ05] Joseph Emerson, Robert Alicki, and Karol Życzkowski. “Scalable noise estimation with random unitary operators”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (Sept. 2005), S347. DOI: 10.1088/1464-4266/7/10/021. URL: <https://dx.doi.org/10.1088/1464-4266/7/10/021>.
- [Ega+21] Laird Egan, Dripto M. Debroy, Crystal Noel, et al. “Fault-tolerant control of an error-corrected qubit”. en. In: *Nature* 598.7880 (Oct. 2021). ISSN: 1476-4687. DOI: 10.1038/s41586-021-03928-y. URL: <https://www.nature.com/articles/s41586-021-03928-y> (visited on 11/17/2021).
- [Eld+18] Zachary Eldredge, Michael Foss-Feig, Jonathan A. Gross, et al. “Optimal and secure measurement protocols for quantum sensor networks”. In: *Phys. Rev. A* 97 (4 Apr. 2018), p. 042337. DOI: 10.1103/PhysRevA.97.042337.
- [EP55] Louis Essen and J. V. L. Parry. “An Atomic Standard of Frequency and Time Interval: A Cæsium Resonator”. In: *Nature* 176.4476 (1955), pp. 280–282.
- [EWP+19] A. Erhard, J.J. Wallman, L. Postler, et al. “Characterizing large-scale quantum computers via cycle benchmarking”. In: *Nature Communications* 10 (2019), p. 5347. DOI: 10.1038/s41467-019-13068-7.
- [Fen+16] Guanru Feng, Joel J Wallman, Brandon Buonacorsi, et al. “Estimating the coherence of noise in quantum control of a solid-state qubit”. In: *Physical review letters* 117.26 (2016), p. 260501.
- [Fey63] Richard P. Feynman. *The Feynman Lectures on Physics*. Vol. 1-3. Addison-Wesley, 1963.
- [Fey82] Richard P Feynman. “Simulating physics with computers”. In: *Int. J. Theor. Phys.* 21.6 (1982), pp. 467–488.
- [FSG09] Austin G. Fowler, Ashley M. Stephens, and Peter Groszkowski. “High-threshold universal quantum computation on the surface code”. In: *Phys. Rev. A* 80 (5 Nov. 2009), p. 052312. DOI: 10.1103/PhysRevA.80.052312.
- [GB18] Riddhi Swaroop Gupta and Michael J. Biercuk. “Machine Learning for Predictive Estimation of Qubit Dynamics Subject to Dephasing”. In: *Phys. Rev. Applied* 9 (6 June 2018), p. 064042. DOI: 10.1103/PhysRevApplied.9.064042.

- [GBC90] Terry Gullion, David B Baker, and Mark S Conradi. “New, compensated Carr–Purcell sequences”. In: *J. Magn. Reson.* 89.3 (1990), pp. 479–484. ISSN: 0022-2364. DOI: [https://doi.org/10.1016/0022-2364\(90\)90331-3](https://doi.org/10.1016/0022-2364(90)90331-3).
- [GC13] Juan Carlos Garcia-Escartin and Pedro Chamorro-Posada. “swap test and Hong-Ou-Mandel effect are equivalent”. In: *Phys. Rev. A* 87 (5 May 2013), p. 052330. DOI: [10.1103/PhysRevA.87.052330](https://doi.org/10.1103/PhysRevA.87.052330). URL: <https://link.aps.org/doi/10.1103/PhysRevA.87.052330>.
- [Gen+14] Genko T. Genov, Daniel Schraft, Thomas Halfmann, et al. “Correction of Arbitrary Field Errors in Population Inversion of Quantum Systems by Universal Composite Pulses”. In: *Phys. Rev. Lett.* 113 (4 July 2014), p. 043001. DOI: [10.1103/PhysRevLett.113.043001](https://doi.org/10.1103/PhysRevLett.113.043001).
- [Giu+20] Tudor Giurgica-Tiron, Yousef Hindy, Ryan LaRose, et al. “Digital zero noise extrapolation for quantum error mitigation”. In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 2020, pp. 306–316. DOI: [10.1109/QCE49297.2020.00045](https://doi.org/10.1109/QCE49297.2020.00045).
- [Got98] Daniel Gottesman. “Theory of fault-tolerant quantum computation”. In: *Phys. Rev. A* 57 (1 Jan. 1998), pp. 127–137. DOI: [10.1103/PhysRevA.57.127](https://doi.org/10.1103/PhysRevA.57.127).
- [Gra+07] M.D. Grace, C. Brif, H. Rabitz, et al. “Fidelity of optimally controlled quantum gates with randomly coupled multiparticle environments”. In: *J. Mod. Opt.* 54.16-17 (2007), pp. 2339–2349.
- [Gra+12] Christopher E. Granade, Christopher Ferrie, Nathan Wiebe, et al. “Robust online Hamiltonian learning”. In: *New J. Phys.* 14 (2012), p. 103013.
- [GS22] Walther Gerlach and Otto Stern. “Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld”. In: *Zeitschrift für Physik* 9 (1922), pp. 349–352.
- [GUB12] Todd Green, Hermann Uys, and Michael J. Biercuk. “High-Order Noise Filtering in Nontrivial Quantum Logic Gates”. In: *Phys. Rev. Lett.* 109 (2 July 2012), p. 020501. DOI: [10.1103/PhysRevLett.109.020501](https://doi.org/10.1103/PhysRevLett.109.020501).
- [Gup+19] Riddhi Swaroop Gupta, Alistair R. Milne, Claire L. Edmunds, et al. “Adaptive scheduling of noise characterization in quantum computers”. In: *Preprint at https://arxiv.org/abs/1904.07225* (2019).
- [Hay+10] D. Hayes, D. N. Matsukevich, P. Maunz, et al. “Entanglement of Atomic Qubits Using an Optical Frequency Comb”. In: *Phys. Rev. Lett.* 104 (14 Apr. 2010), p. 140501. DOI: [10.1103/PhysRevLett.104.140501](https://doi.org/10.1103/PhysRevLett.104.140501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.104.140501>.

- [Her+18] S. Hernández-Gómez, F. Poggiali, P. Cappellaro, et al. “Noise spectroscopy of a quantum-classical environment with a diamond qubit”. In: *Phys. Rev. B* 98 (21 Dec. 2018), p. 214307. DOI: 10.1103/PhysRevB.98.214307.
- [HRB08] H. Häffner, C. F. Roos, and R. Blatt. “Quantum computing with trapped ions”. In: *Phys. Rep.* 469 (2008), p. 155.
- [Joh+15] Blake R. Johnson, Marcus P. da Silva, Colm A. Ryan, et al. “Demonstration of robust quantum gate tomography via randomized benchmarking”. en. In: *New Journal of Physics* 17.11 (Nov. 2015). Publisher: IOP Publishing, p. 113019. ISSN: 1367-2630. DOI: 10.1088/1367-2630/17/11/113019. URL: <https://doi.org/10.1088/1367-2630/17/11/113019> (visited on 03/09/2022).
- [Kab+14] C. Kabytayev, T. J. Green, K. Khodjasteh, et al. “Robustness of composite pulses to time-dependent control noise”. In: *Phys. Rev. A* 90 (2014), p. 012316.
- [Kan+21] Mingyu Kang, Qiyao Liang, Bichen Zhang, et al. “Batch Optimization of Frequency-Modulated Pulses for Robust Two-Qubit Gates in Ion Chains”. In: *Phys. Rev. Appl.* 16 (2 Aug. 2021), p. 024039. DOI: 10.1103/PhysRevApplied.16.024039. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.16.024039>.
- [KB20] Hyeokjea Kwon and Joonwoo Bae. “A hybrid quantum-classical approach to mitigating measurement errors”. In: *arXiv preprint arXiv:2003.12314* (2020).
- [KD18] G.J. Knott and J.A. Doudna. “CRISPR-Cas guides the future of genetic engineering”. In: *Science* 361.6405 (2018), pp. 866–869. DOI: 10.1126/science.aar5903.
- [Kel+14] J. Kelly, R. Barends, B. Campbell, et al. “Optimal Quantum Control Using Randomized Benchmarking”. In: *Phys. Rev. Lett.* 112 (24 June 2014), p. 240504. DOI: 10.1103/PhysRevLett.112.240504.
- [Kim+21] Youngseok Kim, Christopher J. Wood, Theodore J. Yoder, et al. “Scalable error mitigation for noisy quantum circuits produces competitive expectation values”. In: *arXiv:2108.09197 [cond-mat, physics:quant-ph]* (Aug. 2021). arXiv: 2108.09197. URL: <http://arxiv.org/abs/2108.09197> (visited on 11/02/2021).
- [Kit97] A Yu Kitaev. “Quantum computations: algorithms and error correction”. In: *Russ. Math. Surv.* 52 (1997), pp. 1191–1249.
- [KLY15] Shelby Kimmel, Guang Hao Low, and Theodore J. Yoder. “Robust calibration of a universal single-qubit gate set via robust phase estimation”. In: *Phys. Rev. A* 92 (6 Dec. 2015), p. 062315. DOI: 10.1103/PhysRevA.92.062315.

- [KLZ98] E. Knill, R. Laflamme, and Wojciech Zurek. “Resilient quantum computation”. In: *Science* 279 (1998), pp. 342–345.
- [Kni05] Emanuel Knill. “Quantum computing with realistically noisy devices”. In: *Nature* 434 (2005), pp. 39–44.
- [Kra12] Helge Kragh. *Niels Bohr and the quantum atom: the Bohr model of atomic structure, 1913–1925*. Oxford: Oxford University Press, 2012.
- [Kri+21] Sebastian Krinner, Nathan Lacroix, Ants Remm, et al. “Realizing Repeated Quantum Error Correction in a Distance-Three Surface Code”. In: *arXiv:2112.03708 [cond-mat, physics:quant-ph]* (Dec. 2021). arXiv: 2112.03708. URL: <http://arxiv.org/abs/2112.03708> (visited on 01/10/2022).
- [KV09a] Kaveh Khodjasteh and Lorenza Viola. “Dynamical quantum error correction of unitary operations with bounded controls”. In: *Phys. Rev. A* 80 (2009), p. 032314.
- [KV09b] Kaveh Khodjasteh and Lorenza Viola. “Dynamically Error-Corrected Gates for Universal Quantum Computation”. In: *Phys. Rev. Lett.* 102 (8 Feb. 2009), p. 080501. DOI: 10.1103/PhysRevLett.102.080501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.102.080501>.
- [LaR+20] Ryan LaRose, Andrea Mari, Peter J. Karalekas, et al. “Mitiq: A software package for error mitigation on noisy quantum computers”. In: *arXiv:2009.04417 [quant-ph]* (Sept. 2020). arXiv: 2009.04417. URL: <http://arxiv.org/abs/2009.04417> (visited on 01/04/2021).
- [Lee+05] P J Lee, K-A Brickman, L Deslauriers, et al. “Phase control of trapped ion quantum gates”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (Sept. 2005), S371–S383. DOI: 10.1088/1464-4266/7/10/025. URL: <https://doi.org/10.1088/1464-4266/7/10/025>.
- [Leu+18] Pak Hong Leung, Kevin A. Landsman, Caroline Figgatt, et al. “Robust 2-Qubit Gates in a Linear Ion Crystal Using a Frequency-Modulated Driving Force”. In: *Phys. Rev. Lett.* 120 (2 Jan. 2018), p. 020501. DOI: 10.1103/PhysRevLett.120.020501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.020501>.
- [Lev86] M H Levitt. “Composite pulses”. In: *Prog. Nucl. Mag. Res. Spec.* 18.2 (1986), pp. 61–122. ISSN: 00796565. DOI: 10.1016/0079-6565(86)80005-X.
- [Li+19] Muyuan Li, Daniel Miller, Michael Newman, et al. “2D Compass Codes”. In: *Phys. Rev. X* 9 (2 May 2019), p. 021041. DOI: 10.1103/PhysRevX.9.021041.

- [LMP22] Vicente Leyton-Ortega, Swarnadeep Majumder, and Raphael C Pooser. “Quantum error mitigation by hidden inverses protocol in superconducting quantum devices”. In: *Quantum Science and Technology* 8.1 (2022), p. 014008.
- [Lob+21] Daniel Lobser, Joshua Goldberg, Andrew Landahl, et al. *JaqalPaw A Guide to Defining Pulses and Waveforms for Jaqal*. Apr. 2021. URL: <https://qscout.sandia.gov>.
- [LY17] Jing Liu and Haidong Yuan. “Quantum parameter estimation with optimal control”. In: *Phys. Rev. A* 96 (1 July 2017), p. 012117. DOI: 10.1103/PhysRevA.96.012117.
- [LYC14] G. H. Low, T. J. Yoder, and I. L. Chuang. “Optimal Arbitrarily Accurate Composite Pulse Sequences”. In: *Phys. Rev. A* 89 (2014), p. 022341.
- [LYC16] G. H. Low, T. J. Yoder, and I. L. Chuang. “The Methodology of Resonant Equiangular Composite Quantum Gates”. In: *Phys. Rev. X* 6 (2016), p. 041067.
- [Mac+18] Shai Machnes, Elie Assémat, David Tannor, et al. “Tunable, Flexible, and Efficient Optimization of Control Pulses for Practical Qubits”. In: *Phys. Rev. Lett.* 120 (15 Apr. 2018), p. 150401.
- [Maj+22] Swarnadeep Majumder, Christopher G Yale, Titus D Morris, et al. “Characterizing and mitigating coherent errors in a trapped ion quantum processor using hidden inverses”. In: *arXiv preprint arXiv:2205.14225* (2022).
- [Man80] Yu. I. Manin. “Vychislimoe i nevychislimoe [Computable and Noncomputable]”. In: *Sov. Radio* (1980), pp. 13–15.
- [Mau16] Peter Maunz. “High Optical Access Trap 2.0.” In: *Tech. Rep. SAND2016-0796R* (Jan. 2016). DOI: 10.2172/1237003.
- [Mau86] A. A. Maudsley. “Modified Carr–Purcell–Meiboom–Gill sequence for NMR Fourier imaging applications”. In: *J. Magn. Reson* 69 (1986), p. 488.
- [Mav+17] Sandeep Mavadia, Virginia Frey, Jarrah Sastrawan, et al. “Prediction and real-time compensation of qubit decoherence via machine learning”. In: *Nat. Commun.* 8 (Jan. 2017), p. 14106.
- [MB14] T. J. Merrill and K. R. Brown. “Progress in compensating pulse sequences for quantum computation”. In: *Adv. Chem. Phys.* 154 (2014), pp. 241–294.
- [MCB20] Swarnadeep Majumder, Leonardo Andreta de Castro, and Kenneth R. Brown. “Real-time calibration with spectator qubits”. In: *npj Quantum Information* 6 (1 Feb. 2020). DOI: 10.1038/s41534-020-0251-y. URL: <https://doi.org/10.1038/s41534-020-0251-y>.

- [McC+19a] Alex McCaskey, Zachary P. Parks, Jacek Jakowski, et al. “Quantum chemistry as a benchmark for near-term quantum computers”. In: *NPJ Quantum Inf* 5 (2019), p. 99. URL: <https://doi.org/10.1038/s41534-019-0209-0>.
- [McC+19b] Alexander J. McCaskey, Zachary P. Parks, Jacek Jakowski, et al. “Quantum chemistry as a benchmark for near-term quantum computers”. In: *npj Quantum Information* 5.1 (2019), p. 99. URL: <https://doi.org/10.1038/s41534-019-0209-0>.
- [MCF07] S. Montangero, T. Calarco, and R. Fazio. “Robust optimal quantum gates for Josephson charge qubits”. In: *Phys. Rev. Lett.* 99.17 (2007), p. 170501.
- [McK+19] David C. McKay, Sarah Sheldon, John A. Smolin, et al. “Three-Qubit Randomized Benchmarking”. In: *Phys. Rev. Lett.* 122 (20 May 2019), p. 200502. DOI: 10.1103/PhysRevLett.122.200502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.200502>.
- [ME47] John W. Mauchly and J. Presper Eckert. *Electronic numerical integrator and computer*. Patent. U.S. Patent. 1947.
- [Mei+19] Adam M. Meier, Karl A. Burkhardt, Brian J. McMahon, et al. “Testing the robustness of robust phase estimation”. In: *Phys. Rev. A* 100 (5 Nov. 2019), p. 052106. DOI: 10.1103/PhysRevA.100.052106. URL: <https://link.aps.org/doi/10.1103/PhysRevA.100.052106>.
- [MG75] P. Mansfield and P. K. Grannell. ““Diffraction” and microscopy in solids and liquids by NMR”. In: *Phys. Rev. B* 12 (9 Nov. 1975), pp. 3618–3634. DOI: 10.1103/PhysRevB.12.3618. URL: <https://link.aps.org/doi/10.1103/PhysRevB.12.3618>.
- [MGE12] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. “Characterizing quantum gates via randomized benchmarking”. In: *Phys. Rev. A* 85 (4 Apr. 2012), p. 042311. DOI: 10.1103/PhysRevA.85.042311. URL: <https://link.aps.org/doi/10.1103/PhysRevA.85.042311>.
- [Mil+20] Alistair R. Milne, Claire L. Edmunds, Cornelius Hempel, et al. “Phase-Modulated Entangling Gates Robust to Static and Time-Varying Errors”. In: *Phys. Rev. Appl.* 13 (2 Feb. 2020), p. 024022. DOI: 10.1103/PhysRevApplied.13.024022. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.13.024022>.
- [Mon+95] C. Monroe, D. M. Meekhof, B. E. King, et al. “Demonstration of a Fundamental Quantum Logic Gate”. In: *Phys. Rev. Lett.* 75 (25 Dec. 1995), pp. 4714–4717. DOI: 10.1103/PhysRevLett.75.4714. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.75.4714>.

- [Mor+20] B. C. A. Morrison, A. J. Landahl, D. S. Lobser, et al. “Just Another Quantum Assembly Language (Jaql)”. In: *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 2020, pp. 402–408. DOI: 10.1109/QCE49297.2020.00056.
- [Mur+19] Prakash Murali, Jonathan M. Baker, Ali Javadi-Abhari, et al. “Noise-Adaptive Compiler Mappings for Noisy Intermediate-Scale Quantum Computers”. In: *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems. ASPLOS ’19*. Providence, RI, USA: Association for Computing Machinery, 2019, pp. 1015–1029. ISBN: 9781450362405. DOI: 10.1145/3297858.3304075.
- [MYB19] Sam McArdle, Xiao Yuan, and Simon Benjamin. “Error-Mitigated Digital Quantum Simulation”. In: *Phy. Rev. Lett.* 122.18 (2019), p. 180501.
- [MZO20] Filip B Maciejewski, Zoltán Zimborás, and Michał Oszmaniec. “Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography”. In: *Quantum* 4 (2020), p. 257.
- [Nat+21] Paul D. Nation, Hwajung Kang, Neereja Sundaresan, et al. “Scalable mitigation of measurement errors on quantum computers”. In: *arXiv:2108.12518 [quant-ph]* (Aug. 2021). arXiv: 2108.12518. URL: <http://arxiv.org/abs/2108.12518> (visited on 11/02/2021).
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge, UK: Cambridge University Press, 2000.
- [Nie+20] Erik Nielsen, Kenneth Rudinger, Timothy Proctor, et al. “Probing quantum processor performance with pyGSTi”. In: *Quantum Science and Technology* 5.4 (July 2020), p. 044002. DOI: 10.1088/2058-9565/ab8aa4. URL: <https://doi.org/10.1088/2058-9565/ab8aa4>.
- [Nie+21] Erik Nielsen, Kenneth Rudinger, Timothy Proctor, et al. “Efficient flexible characterization of quantum processors with nested error models”. In: *arXiv:2103.02188 [quant-ph]* (Mar. 2021). arXiv: 2103.02188. URL: <http://arxiv.org/abs/2103.02188> (visited on 05/26/2021).
- [NP13] John Napp and John Preskill. “Optimal Bacon–Shor Codes”. In: *Quantum Info. Comput.* 13.5-6 (May 2013), pp. 490–510. ISSN: 1533-7146.
- [Olm+07] S. Olmschenk, K. C. Younge, D. L. Moehring, et al. “Manipulation and detection of a trapped Yb⁺ hyperfine qubit”. In: *Phys. Rev. A* 76 (5 Nov. 2007), p. 052314. DOI: 10.1103/PhysRevA.76.052314. URL: <https://link.aps.org/doi/10.1103/PhysRevA.76.052314>.

- [Piv+21] Christophe Piveteau, David Sutter, Sergey Bravyi, et al. “Error Mitigation for Universal Gates on Encoded Qubits”. In: *Phys. Rev. Lett.* 127 (20 Nov. 2021), p. 200505. DOI: 10.1103/PhysRevLett.127.200505. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.200505>.
- [PLP21] Evan Peters, Andy C. Y. Li, and Gabriel N. Perdue. “Perturbative readout error mitigation for near term quantum computers”. In: *arXiv:2105.08161 [quant-ph]* (May 2021). arXiv: 2105.08161. URL: <http://arxiv.org/abs/2105.08161> (visited on 10/21/2021).
- [PMM07] Line Hjortshøj Pedersen, Niels Martin Møller, and Klaus Mølmer. “Fidelity of quantum operations”. In: *Phys. Lett. A* 367.1 (2007), pp. 47–51. ISSN: 0375-9601. DOI: <https://doi.org/10.1016/j.physleta.2007.02.069>.
- [Pok+18] Bibek Pokharel, Namit Anand, Benjamin Fortman, et al. “Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits”. In: *Phys. Rev. Lett.* 121 (2018), p. 220502.
- [Qia+19] Kevin Qian, Zachary Eldredge, Wenchao Ge, et al. “Heisenberg-scaling measurement protocol for analytic functions with quantum sensor networks”. In: *Phys. Rev. A* 100 (4 Oct. 2019), p. 042304. DOI: 10.1103/PhysRevA.100.042304. URL: <https://link.aps.org/doi/10.1103/PhysRevA.100.042304>.
- [QL13] Gregory Quiroz and Daniel A. Lidar. “Optimized dynamical decoupling via genetic algorithms”. In: *Phys. Rev. A* 88 (2013), p. 052306.
- [Rah64] A. Rahman. “Correlations in the Motion of Atoms in Liquid Argon”. In: *Phys. Rev.* 136 (2A Oct. 1964), A405–A411. DOI: 10.1103/PhysRev.136.A405. URL: <https://link.aps.org/doi/10.1103/PhysRev.136.A405>.
- [RBM18] Nicholas C Rubin, Ryan Babbush, and Jarrod McClean. “Application of fermionic marginal constraints to hybrid quantum algorithms”. In: *New Journal of Physics* 20.5 (May 2018), p. 053020. DOI: 10.1088/1367-2630/aab919. URL: <https://doi.org/10.1088/1367-2630/aab919>.
- [RD10] Chad Rigetti and Michel Devoret. “Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies”. In: *Phys. Rev. B* 81 (13 Apr. 2010), p. 134507. DOI: 10.1103/PhysRevB.81.134507. URL: <https://link.aps.org/doi/10.1103/PhysRevB.81.134507>.

- [RH07] Robert Raussendorf and Jim Harrington. “Fault-Tolerant Quantum Computation with High Threshold in Two Dimensions”. In: *Phys. Rev. Lett.* 98 (19 May 2007), p. 190504. DOI: 10.1103/PhysRevLett.98.190504.
- [RHC10] C. A. Ryan, J. S. Hodges, and D. G. Cory. “Robust Decoupling Techniques to Extend Quantum Coherence in Diamond”. In: *Phys. Rev. Lett.* 105 (20 Nov. 2010), p. 200402. DOI: 10.1103/PhysRevLett.105.200402.
- [RHG07] Robert Raussendorf, Jim Harrington, and K. Goyal. “Topological fault-tolerance in cluster state quantum computation”. In: *N. J. Phys.* 9 (2007), p. 199.
- [Roo08] Christian F Roos. “Ion trap quantum gates with amplitude-modulated laser beams”. In: *New Journal of Physics* 10.1 (Jan. 2008), p. 013002. DOI: 10.1088/1367-2630/10/1/013002. URL: <https://dx.doi.org/10.1088/1367-2630/10/1/013002>.
- [Rud+17] Kenneth Rudinger, Shelby Kimmel, Daniel Lobser, et al. “Experimental Demonstration of a Cheap and Accurate Phase Estimation”. In: *Phys. Rev. Lett.* 118 (2017), p. 190502.
- [Rya+21] C. Ryan-Anderson, J. G. Bohnet, K. Lee, et al. “Realization of Real-Time Fault-Tolerant Quantum Error Correction”. In: *Physical Review X* 11.4 (Dec. 2021). Publisher: American Physical Society, p. 041058. DOI: 10.1103/PhysRevX.11.041058. URL: <https://link.aps.org/doi/10.1103/PhysRevX.11.041058> (visited on 03/30/2022).
- [SÁS12] Alexandre M. Souza, Gonzalo A. Álvarez, and Dieter Suter. “Robust dynamical decoupling”. In: *Phil. Trans. R. Soc. A* 370 (2012), p. 4748.
- [Sch+11] T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, et al. “Optimal control for generating quantum gates in open dissipative systems”. In: *J. Phys. B* 44.15 (2011), p. 154013.
- [Sch26] Erwin Schrödinger. “An Undulatory Theory of the Mechanics of Atoms and Molecules”. In: *Physical Review* 28.6 (1926), pp. 1049–1070.
- [She+16a] Sarah Sheldon, Lev S Bishop, Easwar Magesan, et al. “Characterizing errors on qubit operations via iterative randomized benchmarking”. In: *Physical Review A* 93.1 (2016), p. 012301.
- [She+16b] Sarah Sheldon, Easwar Magesan, Jerry M. Chow, et al. “Procedure for systematically tuning up cross-talk in the cross-resonance gate”. In: *Phys. Rev. A* 93 (6 June 2016), p. 060302. DOI: 10.1103/PhysRevA.93.060302. URL: <https://link.aps.org/doi/10.1103/PhysRevA.93.060302>.

- [Sho94] Peter W. Shor. “Algorithms for quantum computation: Discrete logarithms and factoring”. In: *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (1994), pp. 124–134.
- [Sho95] Peter W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52 (4 Oct. 1995), R2493–R2496. DOI: 10.1103/PhysRevA.52.R2493.
- [Sin+22] Kevin Singh, Conor E Bradley, Shraddha Anand, et al. “Mid-circuit correction of correlated phase errors using an array of spectator qubits”. In: *arXiv preprint arXiv:2208.11716* (2022).
- [SSD16] Pavel Sekatski, Michalis Skotiniotis, and Wolfgang Dür. “Dynamical decoupling leads to improved scaling in noisy quantum metrology”. In: *New J. Phys.* 18.7 (July 2016), p. 073034.
- [Str+20] Armands Strikis, Dayue Qin, Yanzhu Chen, et al. “Learning-based quantum error mitigation”. In: *arXiv preprint arXiv:2005.07601* (2020).
- [Str+21] Armands Strikis, Dayue Qin, Yanzhu Chen, et al. “Learning-Based Quantum Error Mitigation”. In: *PRX Quantum* 2 (4 Nov. 2021), p. 040330. DOI: 10.1103/PRXQuantum.2.040330. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.040330>.
- [Suz+22] Yasunari Suzuki, Suguru Endo, Keisuke Fujii, et al. “Quantum Error Mitigation as a Universal Error Reduction Technique: Applications from the NISQ to the Fault-Tolerant Quantum Computing Eras”. In: *PRX Quantum* 3 (1 Mar. 2022), p. 010345. DOI: 10.1103/PRXQuantum.3.010345. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.3.010345>.
- [TBF18] David K. Tuckett, Stephen D. Bartlett, and Steven T. Flammia. “Ultra-high Error Threshold for Surface Codes with Biased Noise”. In: *Phys. Rev. Lett.* 120 (5 Jan. 2018), p. 050505. DOI: 10.1103/PhysRevLett.120.050505.
- [TBG17] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. “Error Mitigation for Short-Depth Quantum Circuits”. In: *Phys. Rev. Lett.* 119 (18 Nov. 2017), p. 180509. DOI: 10.1103/PhysRevLett.119.180509. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.180509>.
- [Ter15] Barbara M. Terhal. “Quantum error correction for quantum memories”. In: *Rev. Mod. Phys.* 87 (2 Apr. 2015), pp. 307–346. DOI: 10.1103/RevModPhys.87.307.

- [TQ19] Swamit S. Tannu and Moinuddin K. Qureshi. “Not All Qubits Are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers”. In: *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’19. Providence, RI, USA: Association for Computing Machinery, 2019, pp. 987–999. ISBN: 9781450362405. DOI: 10.1145/3297858.3304007.
- [Tuc+19] David K. Tuckett, Andrew S. Darmawan, Christopher T. Chubb, et al. “Tailoring Surface Codes for Highly Biased Noise”. In: *Phys. Rev. X* 9 (4 Nov. 2019), p. 041031. DOI: 10.1103/PhysRevX.9.041031. URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.041031>.
- [TV11] Boyan T. Torosov and Nikolay V. Vitanov. “Smooth composite pulses for high-fidelity quantum information processing”. In: *Phys. Rev. A* 83 (5 May 2011), p. 053420. DOI: 10.1103/PhysRevA.83.053420.
- [Vir+20] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.
- [Vit11] Nikolay V. Vitanov. “Arbitrarily accurate narrowband composite pulse sequences”. In: *Phys. Rev. A* 84 (6 Dec. 2011), p. 065404. DOI: 10.1103/PhysRevA.84.065404.
- [VK03] Lorenza Viola and Emanuel Knill. “Robust Dynamical Decoupling of Quantum Systems with Bounded Controls”. In: *Phys. Rev. Lett.* 90 (3 Jan. 2003), p. 037901. DOI: 10.1103/PhysRevLett.90.037901.
- [VKL99] Lorenza Viola, Emanuel Knill, and Seth Lloyd. “Dynamical Decoupling of Open Quantum Systems”. In: *Phys. Rev. Lett.* 82 (1999), p. 2417.
- [Wal+15] Joel Wallman, Chris Granade, Robin Harper, et al. “Estimating the coherence of noise”. In: *New Journal of Physics* 17.11 (2015), p. 113020.
- [WE16] Joel J. Wallman and Joseph Emerson. “Noise tailoring for scalable quantum computation via randomized compiling”. In: *Physical Review A* 94.5 (Nov. 2016). Publisher: American Physical Society, p. 052325. DOI: 10.1103/PhysRevA.94.052325. URL: <https://link.aps.org/doi/10.1103/PhysRevA.94.052325> (visited on 04/09/2020).
- [Wim94] Stephen Wimperis. “Broadband, narrowband, and passband composite pulses for use in advanced NMR experiments”. In: *J. Magn. Reson.* 109.2 (1994), pp. 221–231. DOI: 10.1006/jmra.1994.1159.
- [Win+92] D. J. Wineland, J. J. Bollinger, W. M. Itano, et al. “Spin squeezing and reduced quantum noise in spectroscopy”. In: *Phys. Rev. A* 46 (11 Dec. 1992), R6797–R6800. DOI: 10.1103/PhysRevA.46.R6797.

- [Win+98] D. J. Wineland, C. Monroe, W. M. Itano, et al. “Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions”. In: *Journal of Research of the National Institute of Standards and Technology* 103 (1998), p. 259.
- [Yet+21] Kübra Yeter-Aydeniz, Bryan T Gard, Jacek Jakowski, et al. “Benchmarking quantum chemistry computations with variational, imaginary time evolution, and Krylov space solver algorithms”. In: *Advanced Quantum Technologies* 4.7 (2021), p. 2100012. URL: <https://doi.org/10.1002/qute.20210001>.
- [ZG20] Alexander Zlokapa and Alexandru Gheorghiu. “A deep learning model for noise prediction on near-term quantum devices”. In: *arXiv preprint arXiv:2005.10811* (2020).
- [Zha+20] Shuaining Zhang, Yao Lu, Kuan Zhang, et al. “Error-Mitigated Quantum Gates Exceeding Physical Fidelities in a Trapped-Ion System”. In: *Nature Communications* 11.1 (Dec. 2020). arXiv: 1905.10135, p. 587. ISSN: 2041-1723. DOI: 10.1038/s41467-020-14376-z. URL: <http://arxiv.org/abs/1905.10135> (visited on 10/01/2020).
- [Zha+22] Bichen Zhang, Swarnadeep Majumder, Pak Hong Leung, et al. “Hidden Inverses: Coherent Error Cancellation at the Circuit Level”. In: *Phys. Rev. Applied* 17 (3 Mar. 2022), p. 034074. DOI: 10.1103/PhysRevApplied.17.034074. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.17.034074>.
- [Zho+18] Sisi Zhou, Mengzhen Zhang, John Preskill, et al. “Achieving the Heisenberg limit in quantum metrology using quantum error correction”. In: *Nat. Commun.* 9 (2018), p. 78.

Biography

Swarnadeep Majumder graduated from Worcester Polytechnic Institute (WPI) with a double major in Physics and Electrical & Computer Engineering in 2018. Afterwards, he joined the Electrical and Computer Engineering PhD program at Duke University, working under the supervision of Dr. Kenneth Brown. During his undergraduate studies, Swarnadeep was awarded the Salisbury Prize and Harold S. Black award from WPI, and he received a National Science Foundation QISE-NET Fellowship during his PhD. Throughout his doctoral studies, he has collaborated with researchers at Oak Ridge National Lab, Sandia National Lab, the University of Chicago, and IBM Quantum. Upon completion of his PhD, Swarnadeep will join IBM Quantum as a researcher.