

DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY


Quantum Machine Learning, Error Correction, and Topological Phases of Matter

Citation

Cong, Iris. 2022. Quantum Machine Learning, Error Correction, and Topological Phases of Matter. Doctoral dissertation, Harvard University Graduate School of Arts and Sciences.

Link

https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37374583

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFlOWI2M2EwYTg

Accessibility

https://accessibility.huit.harvard.edu/digital-accessibility-policy

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story

HARVARD UNIVERSITY
Graduate School of Arts and Sciences



DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Physics

have examined a dissertation entitled

Quantum Machine Learning, Error Correction, and Topological Phases of Matter

presented by Iris Cong

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature _____

Typed name: Professor Mikhail Lukin, Chair

Signature _____

Typed name: Professor Subir Sachdev

Signature Leslie Valiant

Typed name: Professor Leslie Valiant

Date: November 4, 2022

Quantum Machine Learning, Error Correction, and Topological Phases of Matter

A DISSERTATION PRESENTED
BY
IRIS CONG
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
NOVEMBER 2022

©2022 – Iris Cong
all rights reserved.

Thesis advisor: Professor Mikhail D. Lukin Iris Cong

Quantum Machine Learning, Error Correction, and Topological Phases of Matter

Abstract

Quantum computation is an exciting and rapidly developing paradigm of computing, which has the potential to improve or accelerate numerous important applications ranging from cryptography to drug discovery. Inspired by these prospects, researchers and engineers across the world have, in recent years, made tremendous advances in building programmable quantum devices. Nevertheless, the size and error rates in state-of-the-art quantum computers are still several orders of magnitude away from those required to achieve computational speedups using existing quantum algorithms. These developments, along with their limitations, motivate us to consider two important questions: First, are there any useful applications in which current or near-term quantum devices can outperform classical computers? Second, can we develop efficient error-correction schemes to suppress the error rates in near-term quantum devices?

In this dissertation, we present novel solutions to both of these questions. To address the first question, we identify an important class of problems in condensed-matter physics, namely the identification and characterization of quantum phases of matter, as one promising application of current and near-term quantum information processors. We develop new algorithms for using such quantum devices to solve these problems, and we demonstrate how our methods substantially outperform existing approaches. In addition, we respond to the second question by developing hardware-efficient quantum error correction protocols. By leveraging the distinctive properties and

advantages of a given experimental setup to overcome its particular limitations, our proposals significantly reduce the number of quantum bits and operations required for performing quantum error correction and fault-tolerant quantum computation. Together, the works in this dissertation thus pave the way for the near-term realization of quantum computational advantages.

Contents

Title i
Copyright ii
Abstract iii
Table of Contents v
Dedication viii
Acknowledgments ix
Citations to Previous Works xiii

1 Introduction 1
1.1 Background 1
1.2 Key Problems 7
1.3 Main Contributions 8
1.4 Structure of Thesis 9

2 Quantum Phases of Matter 11
2.1 Introduction to Quantum Phases and Phase Transitions 12
2.2 Two-Dimensional Topological Phases of Matter 15
2.3 Symmetry-Protected Topological Phases 18
2.4 Challenge: Quantum Phase Recognition 21

3 Quantum Error Correction 22
3.1 Classical Error Correction 23
3.2 Ideas Behind Quantum Error Correction 25
3.3 Challenge: Fault-Tolerant Quantum Computation 33

4 Quantum Convolutional Neural Networks 35
4.1 Introduction and Motivations 36
4.2 QCNN Circuit Model 37
4.3 Detecting a 1D Symmetry-Protected Topological Phase 41
4.4 Optimizing Quantum Error Correction 48
4.5 Experimental Realizations 51

5 Enhancing Detection of Topological Order by Local Error Correction 52
5.1 Introduction and Motivations . . . 53
5.2 LED Approach . . . 54
5.3 Numerical Detection of Topological Order with Coherent Perturbations . . . 59
5.4 Effect of Incoherent Errors . . . 61
5.5 Arbitrary Local Perturbations . . . 64
5.6 Topological Order Witness . . . 65
5.7 Experimental Realization in Rydberg Atom Arrays . . . 77

6 Hardware-Efficient, Fault-Tolerant Quantum Computation 80
6.1 Introduction and Motivations . . . 81
6.2 Overview of Main Results . . . 86
6.3 Error Channels in Rydberg Atoms . . . 94
6.4 FTQC with the Seven-Qubit Steane Code . . . 101
6.5 Leading-Order Fault-Tolerance with a Repetition Code . . . 116
6.6 Further Considerations Towards Experimental Implementation . . . 129
6.7 Conclusions . . . 133

7 Conclusions and Future Directions 135
7.1 Future Directions for QCNN . . . 136
7.2 Future Directions for LED . . . 137
7.3 Future Directions for Fault-Tolerant Quantum Computation with Rydberg Atoms 138

Appendix A Supplementary Information for Chapter 4 140
A.1 Phase Diagram and QCNN Circuit Simulations . . . 140
A.2 Multiscale String Order Parameters . . . 141
A.3 Construction of QCNN Circuit . . . 142
A.4 QCNN for General QPR Problems . . . 143
A.5 Demonstration of Learning Procedure for QPR . . . 146
A.6 Demonstration of Learning Procedure for QEC . . . 147
A.7 Experimental Resource Analysis . . . 150

Appendix B Supplementary Information for Chapter 5 152
B.1 Numerical Simulations for the Toric Code . . . 152
B.2 Details on Error-Correction and Coarse-Graining Procedures . . . 155
B.3 Patch-Based Decoder . . . 158
B.4 Decoder Dependence . . . 159
B.5 Sample Complexity . . . 161
B.6 Decoder Details for the Ruby Lattice Spin Liquid . . . 162
B.7 Definition and Properties of Fixed-Point States . . . 164

B.8 PEPS Sampling Contraction Details 165
B.9 Calculation of Phase Diagram 167
B.10 Annulus Decoders 169

Appendix C Supplementary Information for Chapter 6 171
C.1 Numerical Computation of Branching Ratios and Transition Rates 171
C.2 An Example of Master Equation Solution for Radiative Decay 174
C.3 Atom Loss Errors 176
C.4 Converting Rydberg Leakage to Pauli Errors 177
C.5 Fault-Tolerant Detection of Rydberg Leakage Errors 179
C.6 Error Syndromes with Postponed Measurements 179
C.7 Implementation of a Bias-Preserving Toffoli Gate 181
C.8 Computing Resource Costs for Rydberg FTQC Protocols 182
C.9 Computing Rydberg Blockade Radius Requirements for Rydberg FTQC Protocols 184
C.10 Blockade Radius Reduction for Ryd-7 187
C.11 Square Lattice Geometry for Ryd-3 189
C.12 Optical Pumping Procedure for the Bias-Preserving CNOT 190

References 205

To all my teachers and mentors,
Especially my first ones, Jing and Jason

Acknowledgments

It takes village to raise a child, and it likewise took a community for me to pursue and attain my PhD. Over the past five years, I have been exceptionally fortunate and privileged to have the mentorship, collaborations, and support of so many incredible people; I now take this opportunity to express my gratitude for some of the most vital relationships that enabled this work.

My PhD advisor, Prof. Mikhail (Misha) Lukin, has played a pivotal role in overseeing and enabling my growth from a student to a researcher. I am always amazed and inspired by how he draws incredible connections between seemingly unrelated concepts; indeed, my collaborators and I often initially dismiss these connections as impossible, only to re-discover them several months later. Moreover, Misha has asked so many insightful questions, prompting me to consider additional directions or applications. This frequently came to my chagrin when I thought I was nearly done with a project, but ultimately improved my work consistently. Finally, I appreciate how he allowed and encouraged me to be myself and contribute uniquely to the research group. I will always be grateful for how Misha thus helped me grow into an interdisciplinary researcher.

I also thank the other members of my committee, Profs. Subir Sachdev and Leslie Valiant, for their feedback on my work and their guidance in my development as a scientist. In addition, I thank three other professors who advised me during these years, Profs. Eugene Demler, Susanne Yelin, and Hannes Pichler.

Next, I thank the mentors and collaborators I had on each of my PhD projects. During my first and second years, I was privileged to work with Soonwon Choi. As I had studied computer sci-

ence (CS) in my undergraduate studies, Soonwon taught me so much of the physics of quantum information that I know and use. I am also especially grateful to Shengtao Wang, who oversaw and mentored me throughout my second project. Finally, I would like to particularly thank Nishad Maskara: given our diverse and complementary skill sets, our collaboration was remarkably exciting, productive, and essential to the success of my final PhD project.

In addition to my theory collaborators, I would also like to give a special shout-out to the members of the atom array experimental team, and in particular my co-authors Harry Levine, Alexander Keesling, Dolev Bluvstein, and Giulia Semeghini. They all taught me so much about the experimental physics behind Rydberg-atom systems—the hardware platform which I consider throughout this thesis—and they played an incredible role of making the dreams of us theorists come true.

Throughout my PhD, I have also had the pleasure to work with and learn from various other talented researchers, including Minh Tran, Thomas Schuster, Bryce Kobrin, Emil Khabiboulline, Sona Najafi, Giuliano Giudici, Prof. Johannes Borregaard, and Prof. Norman Yao. In addition, I have been privileged to interact with many brilliant members of the Lukin group, including Maddie Cain, Sambuddha Chattopadhyay, Joonhee Choi, Tamara Dordevic, Dirk Englund, Ruffin Evans, Xun Gao, Nate Gemelke, Wen Wei Ho, Marcin Kalinowski, Erik Knall, Renate Landig, Nathaniel Leitao, Leigh Martin, Maya Miklos, Paloma Ocola, Kristine Rezai, Ralf Riedinger, Emma Rosenfeld, Polnop Samutpraphoot, DaLi Schaefer, Chi Shu, Pavel Stroganov, Tamara Sumarac, Andrey Sushko, Aziza Suleymanzade, Nazlı Uğur Köylüoğlu, Katherine Van Kirk, Leo Zhou, and Hengyun Zhou. Furthermore, I deeply appreciate several friendships I made with other Harvard students, such as Angela Ma, Yanting Teng and Maria Tikhanovskaya.

Moreover, this work would not have been possible without the support of many staff members in the physics department: Samantha Dakoulas, Lisa Cacciabaudo, Carol Davis, and Prof. Jacob Barandes. I am particularly indebted to Jacob: from my first month as a PhD student, he crucially helped me transition from CS to physics and make up missed curriculum. When I later faced var-

ious academic and personal challenges, he was a guiding light who empowered me to regain my courage and confidence and discover my unique identity, so that I could carry on and follow my dreams. I truly do not know where I would be today without his support and guidance.

Furthermore, I thank the three graduate research fellowships that funded my studies over the last five years: the National Defense Science and Engineering Graduate (NDSEG) Fellowship, the Paul and Daisy (PD) Soros Fellowship, and the Fannie and John Hertz Fellowship. In particular, the PD Soros and Hertz communities helped me broaden my horizons and build particularly meaningful connections outside my immediate field of study. I especially thank Drs. Alfred Spector and Rhonda Kost, who funded my Hertz Fellowship; I enjoyed and learned from many conversations we had about my research and career goals.

I now acknowledge several people who played pivotal roles in helping me enter the field of quantum information. The vision of Dr. Norman Mirman and Mrs. Beverly Mirman, together with the incredible teaching and inspiration of Mr. Larry Wiener, were indispensable for me to discover my talent and passion for math and science. I am also indebted to the late Prof. Michael Jura, my first physics professor; I truly wish I could have told him how a morning conversation we had led me to this research field. Finally, I thank my undergraduate advisors, Profs. Luming Duan, Zhenghan Wang, and Michael Freedman for everything they did to prepare and inspire me to pursue this PhD.

During my time in graduate school, I experienced several medical challenges; I therefore thank four phenomenal physicians without whom this thesis work would not be possible: Drs. Christos D. Photopoulos, Arun J. Ramappa, Ariel Seroussi, and Mary F. Youssef.

Throughout my PhD studies, I have been particularly fortunate to have several incredible friends who created so much happiness in my life, and who provided indispensable compassion and strength to uplift me when I faced my greatest challenges. To Sarah Battat, Luke Harmon, Maya Miklos, Katherine Van Kirk, and Grace Zhang—thank you for your continual love and support, and for all the extraordinary experiences we have shared over the years. In addition, I am grateful to Ms. Yuliya

Barsky, who introduced me to the wonders of music and has known me since I was four years old, and who continues to enrich my scientifically-focused life with such beautiful, musical and artistic adventures.

Lastly, my deepest thanks go to my family. My grandparents, Manzhi and Keli, have brought me so much joy ever since my childhood and youth, and they are two of the most wonderful cheerleaders I have. And to my parents, Jing and Jason—no words are adequate to express my gratitude for all the love and care you have given me over the last 23 years, for the sacrifices you have made in raising me, and for your unwavering support throughout all my highs and lows. Without you, I could never have become the person and scholar that I am today.

Citations to Previous Works

The following chapters of this thesis were co-authored and previously published:

Chapter 4 and Appendix A were reproduced from:

> Iris Cong, Soonwon Choi, and Mikhail D. Lukin, Quantum Convolutional Neural Networks, *Nature Physics* **15**, 1273-1278 (2019).

Chapter 5 and Appendix B were reproduced from:

> Iris Cong*, Nishad Maskara*, Minh C. Tran, Hannes Pichler, Giulia Semeghini, Susanne F. Yelin, Soonwon Choi, and Mikhail D. Lukin, Enhancing Detection of Topological Order by Local Error Correction, *arXiv*:2209.12428 (2022).

Chapter 6 and Appendix C were reproduced from:

> Iris Cong, Harry Levine, Alexander Keesling, Dolev Bluvstein, Sheng-Tao Wang, and Mikhail D. Lukin, Hardware-Efficient, Fault-Tolerant Quantum Computation with Rydberg Atoms, *Physical Review X* **12**, 021049 (2022).

Furthermore, portions of Chapter 3 of this thesis are taken, with permission, from course notes which I typeset for Prof. Mikhail Lukin's graduate class on quantum information (Physics 271).

*These authors contributed equally to this work.

Non est ad astra mollis e terris via.

Lucius Annaeus Seneca

1 Introduction

1.1 Background

Throughout the last half-century, rapid progress and developments in technology have revolutionized our world and our daily lives. Notable examples include the development of personal computers and cellular communication, or of big data and machine learning techniques. Many of these breakthroughs can be accredited to the so-called Moore's Law scaling[119], which projects the dou-

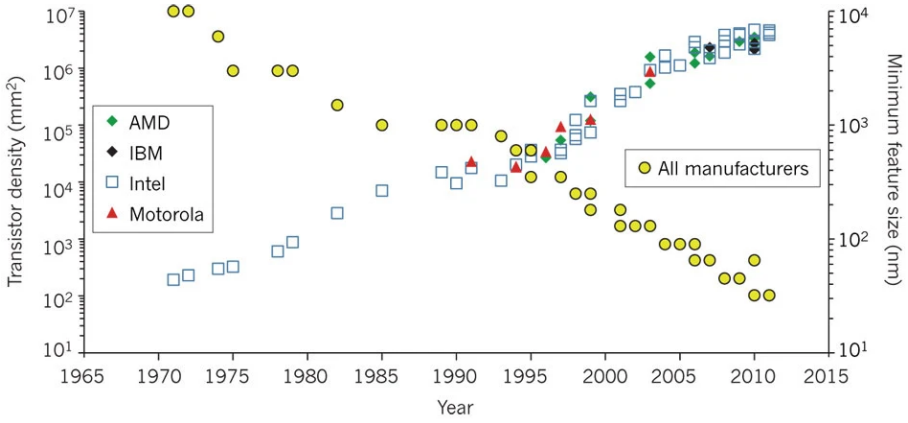


Figure 1.1: Moore's Law scaling. The size of individual transistors has decreased exponentially over time in the last several decades. Since these transistors form the building blocks of modern electronics, this has resulted in an exponential increase in the computational power of microprocessors or cores in devices such as our personal computers or cell phones. Figure credit to Ref. [56].

bling every two years in the number of components (transistors) in an integrated circuit, such as the microprocessors or cores used in our personal computers and cell phones. Over the past several decades, this empirical prediction has guided and motivated the semiconductor industry to develop novel methods to reduce transistor sizes and thereby increase the transistor density and computational power per core. Indeed, while transistors sizes were on the order of one micrometer in 1960, they have now reduced to several nanometers (see Refs. [56,113] and Figure 1.1). Unfortunately, this exponential scaling must inevitably come to an end: a 2 nm transistor, such as the one recently announced by IBM [113], has a width which is only 10 to 20 times the diameter of each silicon atom within the transistor. At such length-scales, quantum-mechanical effects arise; because such effects are not considered in the classical, digital model of computation, new computing paradigms are required to continue the consistent increase in computational power that enables developments of newer, more complex technologies such as faster smartphones or self-driving cars.

One promising solution to this upcoming challenge is quantum computing, a rapidly developing, exciting computational paradigm with the potential to accelerate numerous applications

ranging from cryptography to drug discovery. Indeed, this paradigm explicitly leverages quantum-mechanical properties to achieve computational speedup, and its individual components can be as small as a single atom. In this model of computation, information is encoded in *qudits*, which are quantum systems that can exist in d possible states, $|0\rangle, |1\rangle, \ldots, |d-1\rangle$. For the present discussion and throughout much of this thesis, we focus on the typical choice of $d = 2$ (where the qudits are called *qubits*); this could, for example, be realized through the spin of an electron which can be oriented either upwards or downwards. In certain later sections, we will also consider the more general scenario and specify this explicitly.

Quantum mechanically, the state of a single qubit can be represented as a generic superposition $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where the coefficients α and β are complex numbers satisfying a normalization condition $|\alpha|^2 + |\beta|^2 = 1$. More generally, a string of n qubits can be described as a superposition of 2^n possible states:

$$|\psi^n\rangle = \sum_{b_0\in\{0,1\}} \sum_{b_1\in\{0,1\}} \ldots \sum_{b_{n-1}\in\{0,1\}} \alpha_{b_1,b_2,\ldots b_{n-1}} |b_0 b_1 \ldots b_{n-1}\rangle \tag{1.1}$$

Here, the coefficients are again complex numbers satisfying

$$\sum_{b_0\in\{0,1\}} \sum_{b_1\in\{0,1\}} \ldots \sum_{b_{n-1}\in\{0,1\}} \left|\alpha_{b_1,b_2,\ldots b_{n-1}}\right|^2 = 1.$$

We notice two important distinctions between the quantum and classical models of computation. First, while the basic unit of classical information, a bit, can only be in one of two states, 0 or 1, a qubit can be in a superposition of the two quantum states $|0\rangle$ and $|1\rangle$. Second, and more importantly, a classical string of n bits can be described in terms of the state of each bit individually; this is not possible for a generic n-qubit quantum state of the form shown in Equation (1.1), as the state of any one qubit can depend on the state of the other $n-1$ qubits. This second property of quantum

states is known as *quantum entanglement*, and it and forms the basis of the computational speedup obtained in quantum computing. In particular, while a classical string of n bits can be describing by using only n numbers, specifying an n-qubit quantum state in general requires 2^n complex numbers—an enormous amount of information even for modest values of $n \sim 50 - 100$. These 2^n complex numbers can be represented as a vector which belongs to a 2^n-dimensional Hilbert space, which is the n-qubit Hilbert space. Thus, the presence of quantum entanglement makes classical simulation of generic quantum systems particularly difficult and is believed to differentiate the computation power of classical and quantum computers.

Given a set of n qubits, the operations that can be performed on these qubits are known as *quantum gates* or *unitary gates*. The most general quantum gate acting on n qubits can be represented a $2^n \times 2^n$ unitary transformation acting on the n-qubit Hilbert space. However, most $2^n \times 2^n$ unitary gates are extremely difficult to implement experimentally, as they require high-fidelity (i.e., low-noise) manipulation and entanglement of all n qubits simultaneously. Instead, these complex operations are typically broken down into a series of single-, two-, or three-qubit gates, which act nontrivially on one, two, or three out of the n qubits, respectively, while leaving the remaining qubits unchanged. Two examples of such single-qubit gates are the Pauli-X and Pauli-Z gates, which perform bit-flips or phase-flips on a single qubit:

$$X : |0\rangle \leftrightarrow |1\rangle, \qquad Z : |+\rangle \leftrightarrow |-\rangle \tag{1.2}$$

where $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$ denote superposition states. Likewise, we can consider two examples of

two-qubit gates:

$$\mathrm{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \qquad \mathrm{CZ} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}. \tag{1.3}$$

The CNOT gate performs a bit-flip on the second (a.k.a. "target") qubit conditioned on the state of the first (a.k.a. "control") qubit, while the CZ gate performs a phase-flip on the target qubit conditioned on the state of the control qubit. The decomposition of an n-qubit unitary operation into d-qubit unitaries (for small d, such as 1, 2, or 3) is known as a *quantum circuit* realization of the n-qubit operation.

While quantum circuit decompositions are already very helpful for efficient and high-fidelity implementation, many experiments impose two further conditions, namely the *locality* of multi-qubit gates and the *depth* of the overall quantum circuit. More specifically, given any geometric arrangement pattern of qubits (e.g., a square lattice), a multi-qubit gate is said to be local if the distance between any two qubits involved the gate can be bounded by some small, finite value which remains constant as the number n of qubits is increased. Local gates are typically easier to implement than non-local ones, as the involved qubits are closer together, which facilitates high-fidelity entangling interactions between these qubits. One further useful property of local gates involving few qubits is that a sequence of such gates can be performed in parallel, if the gates act on disjoint sets of qubits. Upon applying all such parallelizations in a given quantum circuit, one can then define the depth of a quantum circuit to be the length of the longest sequence of gates involving a single qubit, which cannot be performed in parallel with one another. Increasing the depth of a quantum circuit often increases the complexity of the n-qubit transformation it implements, but it typically also increases

the difficulty of the circuit's experimental implementation. As such, one class of circuits which is frequently studied is that of finite-depth local unitary circuits—namely, quantum circuits whose depths are bounded by some fixed constant d which remains constant as the number of qubits n is increased. This class of circuits becomes particularly useful when studying quantum phases of matter, as seen in subsequent chapters of this thesis.

Given a quantum computer with a large number n of qubits on which operations can be performed with very high fidelity, quantum algorithms have been proposed for solving several important tasks with provable speedup compared to the best-known classical algorithms. For instance, Shor's algorithm for factoring integers [145] is believed to have exponential speedup over any classical factoring algorithm. This has momentous implications for cryptography, as it could be used to efficiently break cryptographic schemes which are widely used for encrypting email and other digital transactions performed on the Internet [132]. Another notable quantum algorithm with provable speedup is Grover's algorithm for searching through an unstructured database, which has a quadratic advantage in the number of database entries when compared to classical methodologies [67].

Motivated by these exciting possibilities, many groups across the globe have made extraordinary advances in experimental quantum information, especially during this past decade. Moreover, multiple potential hardware platforms have emerged for realizing quantum computation, including neutral Rydberg atoms, superconducting circuits, trapped ions, among many others [44,27,72,95,136]. Given this remarkable progress, it is reasonable to envision the development of quantum computers with $\sim 10,000 - 100,000$ qubits and gate fidelities of $\sim 99.999\%$ within the next decade. Such systems are likely impossible to simulate classically. However, they are not yet sufficient for observing the true quantum advantage by applying the aforementioned impactful algorithms.

1.2 Key Problems

The above developments and considerations motivate us to ask two important questions regarding how to use and improve state-of-the-art quantum computers:

1. Given the apparent difficulty of classically simulating large quantum systems, are there any useful applications in which current or near-term quantum devices can outperform classical computers?

2. Can we develop efficient error-correction schemes to suppress the error rates and further improve qubit and gate fidelities in near-term quantum devices?

While approaches have been considered throughout the literature for addressing each of these questions, fully practical and useful solutions still remain elusive. For example, two groups have recently demonstrated "quantum supremacy" by using state-of-the-art quantum computers to sample from a specific probability distribution associated with a pseudo-random quantum circuit. In these prominent experiments, the quantum computers were observed to perform significantly faster than the projected runtime for any known classical algorithm, even when using the world's fastest supercomputers[3,183]. Nevertheless, these works only partially address Question 1 above, as it is highly unclear whether there is any practical usage for the algorithms that were performed.

To answer Question 2, many elegant, theoretical proposals have been developed for performing quantum error correction, where the state of an individual, *logical* qubit is delocalized and encoded using multiple physical qubits (see Chapter 3 and Ref. [155]). The redundant encoding then ensures that the state of the logical qubit can be recovered, even when a small number of the underlying physical qubits are damaged by noise. However, due to the complexity of such encodings, it becomes particularly difficult to perform error-corrected, fault-tolerant operations on the logical qubits: indeed, performing certain logical operations may then require hundreds of physical qubits

and/or thousands of physical gates, which is highly impractical for near-term quantum devices (see Chapter 6). The problem of developing efficient quantum error correction schemes thus remains important and largely unsolved.

1.3 Main Contributions

In this thesis, we present novel solutions to both of the questions above. As one candidate answer to Question 1, we consider an important problem in condensed-matter and quantum physics—namely, the recognition and characterization of quantum phases of matter. To solve this problem, we first design a new, machine-learning-inspired quantum algorithm, the quantum convolutional neural network (QCNN), which is realizable on near-term quantum devices. We find that QCNN substantially outperforms existing approaches for recognizing and characterizing one-dimensional quantum phases. Moreover, we find that this exceptional performance of our algorithm can be understood through a combination of well-known theoretical concepts ranging from condensed-matter and high-energy physics, to computer science and quantum information.

Motivated by this success, we further extend this theoretical framework to detect and characterize two-dimensional, topologically ordered states of matter, which have important implications for many subfields of physical sciences and engineering. Our approach, locally error-corrected decoration (LED), allows for particularly efficient and robust identification of topological order, especially when compared to existing approaches. We demonstrate the power of LED through numerical simulations of Kitaev's toric code model[92], a canonical example of topological order, and we subsequently apply it to a recent experimental realization of a quantum spin liquid using a Rydberg-atom quantum simulator[143].

In response to Question 2, we notice that the exceedingly large cost of performing fault-tolerant quantum computation can be largely accredited to the fact that most existing approaches attempt to

protect against a generic set of errors which can be present in any hardware platform for quantum computation, rather than targeting the specific errors in the experimental platform under consideration. As such, we propose *hardware-efficient* methods for quantum error correction and fault-tolerant quantum computation, which are designed to best mitigate errors in the quantum device. As a first example, we find that the same QCNN algorithm with which we study quantum phases of matter can also be used to optimize a logical qubit encoding scheme based on the underlying error model of any given experiment. In a subsequent work, we focus specifically on the class of quantum computers based on arrays of neutral atoms, where qubits are encoded in non-interacting ground states of neutral alkali or alkaline earth(-like) atoms, and multi-qubit quantum gates are performed by exciting the atoms to highly-excited states (the so-called "Rydberg states"). In particular, we provide the first complete characterization of the dominant sources of error in such a Rydberg-atom quantum computer. We then leverage the unique capabilities of neutral atoms to design hardware-efficient, fault-tolerant quantum computation schemes that mitigate these errors. Notably, we develop a novel and distinctly efficient method to address the most important errors, which are *leakage errors* associated with the decay of atomic qubits to states outside of the computational subspace. These advances allow us to significantly reduce the resource cost for fault-tolerant quantum computation compared to existing, general-purpose schemes.

1.4 Structure of Thesis

The rest of this thesis is structured as follows. Chapters 2 and 3 provide background on quantum phases of matter and quantum error correction, respectively; this forms the foundation on which our results are built. In Chapter 4, we present a new quantum machine-learning based paradigm, the quantum convolutional neural network (QCNN). We show that QCNNs are powerful both for recognizing exotic quantum phases of matter and for efficiently optimizing logical qubit encodings

depending on which errors are present in the experimental system at hand. We then delve deeper into each of these questions: In Chapter 5, we design a new framework which is inspired by QCNN and error-correction techniques, namely locally error-corrected dressing (LED), and demonstrate its success at detecting and characterizing two-dimensional topological phases of matter. Subsequently, in Chapter 6 we present our novel approaches for performing hardware-efficient, fault-tolerant quantum computation on Rydberg-atom platforms, and we illustrate the striking advantages of these approaches compared to existing, general-purpose methods. Finally, we conclude in Chapter 7 with a discussion of the exciting future directions for each of these projects.

The more that you read, the more things you will know.

The more that you learn, the more places you'll go.

Dr. Seuss

2

Quantum Phases of Matter

With many scientific and practical implications, the study of different phases of matter constitutes an important focus of theoretical and experimental physics. For example, classical phase transitions, such as transitions between the solid and liquid phases of matter, account for many important observations in our everyday lives. Meanwhile, in the quantum-mechanical realm, the theory of quantum phases transitions elucidates scientifically and technologically significant phe-

nomena such as the superconductor-insulator and metal-insulator transitions[6,135]. Such discoveries have led to a widespread interest in the subject of quantum phases and phase transitions.

In this chapter, we provide a brief introduction to this subject. We subsequently discuss a particular class of quantum phases—the so-called topological and symmetry-protected topological phases—which have recently attracted attention due to their technological implications, such as their potential to aid in the realization of scalable, fault-tolerant quantum computation. This then leads us to pose an important problem which we study throughout this thesis: how to efficiently and accurately identify and characterize different quantum phases of matter.

2.1 Introduction to Quantum Phases and Phase Transitions

Quantum phases and phase transitions describe the behavior of many-body systems at zero temperature, governed by quantum-mechanical effects. In the simplest example, we can consider a family of Hamiltonians $H(g)$ which are characterized by a single, dimensionless coupling constant g. We suppose that each $H(g)$ acts on the same degrees of freedom, which reside on the sites of a lattice. When this underlying lattice is finite, the ground state energy is generically a smooth, analytic function of g; however, in the case of an infinite lattice, non-analyticities can emerge at specific *critical points* where the energy gap between a ground and excited state vanishes. Such a non-analyticity then corresponds to a *quantum phase transition*: typically, it results in a significant qualitative change in the properties of the ground state of $H(g)$ [135]. Additionally, as we will see throughout this chapter, the ability to adiabatically deform one Hamiltonian (e.g., $H(g_1)$ for some g_1) to another (e.g., $H(g_2)$) without closing an energy gap plays an important role in determining whether their ground states belong to the same quantum phase of matter.

1D Quantum Ising Model

As a concrete example, we consider the 1D quantum Ising model, where qubits are placed on along a one-dimensional chain. We begin by defining the *Pauli matrices* acting on any given qubit:

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \qquad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{2.1}$$

The quantum Ising Hamiltonian $H(g)$ is then expressed in terms of Pauli matrices acting on each qubit:

$$H_I(g) = -J \sum_i Z_i Z_{i+1} - Jg \sum_i X_i. \tag{2.2}$$

When $g = 0$, H_I reduces to the familiar classical Ising Hamiltonian; however, for nonzero g, the second term in (2.2) induces quantum-mechanical fluctuations which can flip qubit values.

To study the ground state characteristics given a generic value of g, we can consider two limiting cases. First, when $g \gg 1$, the second term dominates, so the ground state $|\psi_0\rangle$ satisfies $X_i|\psi_0\rangle = |\psi_0\rangle$ for all i and is given by $|\psi_0\rangle = \otimes_i |+\rangle_i$, where $\otimes$ denotes the tensor product, and

$$|\pm\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle \pm |1\rangle \right). \tag{2.3}$$

On the other hand, when $g \ll 1$, the first term dominates; indeed, in the classical case of $g = 0$, each spin must be in an eigenstate of the Pauli-Z operator ($|0\rangle$ or $|1\rangle$). In this case, because the operators $Z_i Z_{i+1}$ are mutually commuting for all i, they can be simultaneously diagonalized, so that the ground state lies in the $+1$ eigenspace of each such operator. In particular, there are two ground states in this case, namely $|0\rangle = \otimes_i |0\rangle_i$ and $|1\rangle = \otimes_i |1\rangle_i$. More generally, one can show that this two-fold ground state degeneracy survives for any small $g \ll 1$. This is because the Hamiltonian H_I

has an exact, global $\mathbb{Z}_2$ symmetry transformation

$$Z_i \mapsto -Z_i, \qquad X_i \mapsto X_i, \tag{2.4}$$

as H_I remains invariant upon applying the transformation. The symmetry is generated by the unitary operator $\prod_i X_i$, and it interchanges the two ground states $|0\rangle$ and $|1\rangle$.

As such, when any thermodynamic system of $H(g \ll 1)$ is cooled to zero temperature, it must "choose" one of the ground states; this is the so-called "spontaneous symmetry breaking" of the $\mathbb{Z}_2$ symmetry. In contrast, $H(g)$ has only a single ground state at large g, which is invariant under the $\mathbb{Z}_2$ symmetry transformation. These considerations suggest the presence of two distinct quantum phases of matter in this model, which are separated by a critical point g_c: for $g > g_c$, the ground state belongs to a *paramagnetic* phase (sometimes called a *disordered* phase), while for $g < g_c$, the ground state belongs to a *magnetically ordered* phase (also referred to as the spontaneous symmetry-broken phase). Indeed, one finds that paramagnetic states cannot be transformed into magnetically ordered states (and vice versa) via a finite-time adiabatic Hamiltonian evolution; furthermore, the two cases can be distinguished by measuring the expectation value $\langle Z_i \rangle$ on any site i, a local *order parameter* which is nonzero in the magnetically ordered phase and zero throughout the paramagnetic phase.

As seen above, the two phases in the quantum Ising model can be characterized by the action of the $\mathbb{Z}_2$ symmetry on a ground state of $H(g)$. More generally, symmetry transformations and the associated symmetry-breaking can be used to characterize a plethora of quantum phases of matter. This led to a general theory of quantum phases and phase transitions, developed by Landau [96], where different phases were differentiated based on their symmetries. Later, Ginzburg and Landau [62] introduced generic order parameters which transform nontrivially under the symmetry transformations to detect the corresponding phases. In this formalism, every quantum phase of matter can be characterized by a pair of groups $G_\psi \subseteq G_H$, where G_H is the symmetry group of the system's

Hamiltonian, and G_ψ is the unbroken symmetry group of the equilibrium state.

This elegant theory of quantum phases and phase transitions has been largely successful: it allows one to classify all 230 different kinds of three-dimensional crystals and compute important critical properties associated with phase transitions, and can be used to understand superconducting and superfluid states. Nevertheless, theoretical and experimental developments in the last few decades have shown that this theory is incomplete: for example, different fractional quantum Hall states, which belong to different phases of matter, all have the same symmetry [152,97,169]. Such observations have led to paradigms for understanding these new, more exotic phases of matter, several of which are discussed in the rest of this chapter.

2.2 Two-Dimensional Topological Phases of Matter

To overcome the limitations of Landau theory, we return to the understanding of quantum phase transitions in terms of the energy gap between ground and excited states of the Hamiltonians under consideration. Using this understanding, one can define a (gapped) quantum phase of matter as an equivalence class of gapped Hamiltonians—Hamiltonians that have a finite energy gap between the ground and excited states, independent of the system size—which can be analytically deformed into one another without closing the energy gap [169]. Equivalently, one may also view this equivalence relation in terms of the ground states of the relevant Hamiltonians, where two ground states $|\psi_1\rangle$ and $|\psi_2\rangle$ belong to the same phase if and only if they are related by a local unitary transformation—that is, if one can transform $|\psi_1\rangle$ into $|\psi_2\rangle$ by applying a finite-depth sequence of local quantum gates (see Chapter 1).

In general, local quantum operations can change the short-range entanglement structure of quantum states, namely the entanglement properties between qubits separated by a finite distance that is independent of the system size. Meanwhile, they cannot change the long-range entanglement

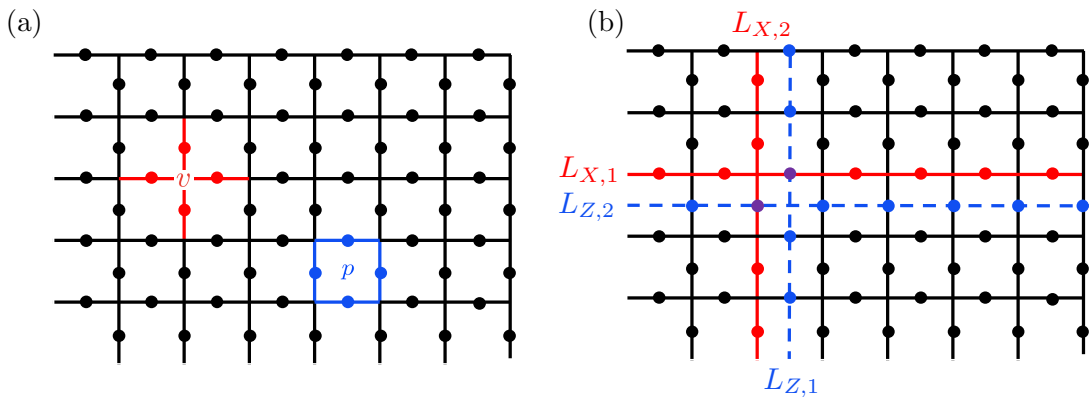


Figure 2.1: Kitaev's toric code model. (a) Qubits are placed on the edges of a square lattice, and the Hamiltonian consists of vertex and plaquette terms as highlighted. (b) When the lattice is placed on a torus, the top and bottom boundaries in this illustration are identified, and then the left and right boundaries are identified. The ground state space in this case is four-fold degenerate and can be used to encode two qubits. The logical operators for these qubits are supported on axial loops (e.g., $L_{X,2}$ and $L_{Z,1}$) and equatorial loops (e.g., $L_{X,1}$ and $L_{Z,2}$) of the torus as illustrated.

structure between regions of qubits separated by an extensive distance. As a result, the above definition distinguishes different quantum phases based on their long-range entanglement structure. In this setting, if a state can be connected via a finite-depth local unitary transformation to a trivial product state, such as a state with all qubits in the $|0\rangle$ state, it belongs to a trivial phase with no long-range entanglement. On the other hand, any other state must exhibit long-range entanglement and is considered to be *topologically ordered*; furthermore, the different equivalence classes characterize different kinds of topological order. This definition of quantum phases can therefore distinguish two systems with the same symmetries if they exhibit different long-range entanglement structures, thereby circumventing some of the limitations of Landau's theory.

Kitaev's Toric Code Model

As a concrete example of a topological phase of matter, we consider the paradigmatic toric code model introduced by Kitaev[92]. In this model, qubits are localized on the edges of a square lattice,

and the Hamiltonian is given by

$$H_{\mathrm{TC}} = -J\sum_{v} A_v - J\sum_{p} B_p \tag{2.5}$$

where

$$A_v = \prod_{i\in \mathrm{adj}(v)} X_i, \qquad B_p = \prod_{i\in \mathrm{adj}(p)} Z_i \tag{2.6}$$

are *stabilizer operators* of the toric code, and $\mathrm{adj}(v)$ and $\mathrm{adj}(p)$ denote the set of edges touching a given vertex v or forming a plaquette p of the lattice, respectively (Figure 2.1(a)). Notice how all terms A_v and B_p are mutually commuting, so that they can be simultaneously diagonalized. This implies that any ground state of H_{TC} must satisfy $A_v = 1$ at every vertex, and $B_p = 1$ at every plaquette.

To see how H_{TC} gives rise to topological ordering, we consider different scenarios in which the square lattice is placed on surfaces with differing topology. For a spherical surface, the ground state of H_{TC} is unique, as there is only one state satisfying $A_v = 1$ at every vertex and $B_p = 1$ at every plaquette. On the other hand, when the square lattice is placed on a torus, the ground state space is four-fold degenerate: let $L_{X,1}$ and $L_{X,2}$ be particular choices of non-contractible axial and equatorial loops of the torus, respectively, supported on the edges of our lattice. Furthermore, let $L_{Z,1}$ and $L_{Z,2}$ be the corresponding axial and equatorial loops supported on the dual lattice, which is defined as the lattice where the vertices and plaquettes are switched (see Figure 2.1(b)). Then one can define logical operators $X_{L_1} = \otimes_{i\in L_{X,1}} X_i$, $X_{L_2} = \otimes_{i\in L_{X,2}} X_i$, $Z_{L_1} = \otimes_{i\in L_{Z,1}} Z_i$, and $Z_{L_2} = \otimes_{i\in L_{Z,2}} Z_i$. Notice how, in this case, each of the four operators $X_{L_1}, X_{L_2}, Z_{L_1}$, and Z_{L_2} commutes with all the stabilizer operators (2.6), but cannot be written as any product of stabilizers. However, X_{L_1} anticommutes with Z_{L_2}, and X_{L_2} anticommutes with Z_{L_1}. This implies that H_{TC} must have four degenerate ground states on a torus, which span the simultaneous $+1$ eigenspace of all stabilizers.

Moreover, this ground state space forms a quantum error-correcting code (see Chapter 3): all local operators either act trivially on ground states or couple them to excited states[92]. This means that the topological ordering and associated ground state degeneracy is robust to small, local perturbations of the Hamiltonian H_{TC}. In particular, to convert between any two distinct ground states of H_{TC}, one must apply an operator with support comparable to the lattice height or width.

Due to the changes in ground state degeneracy upon varying the topology of the underlying manifold, H_{TC} (and locally perturbed variants) must belong to a different phase of matter than a trivially-ordered Hamiltonian such as $H_X = \otimes_i X_i$, which has a single ground state (a trivial product state) independent of topology. However, the toric code Hamiltonian and ground states have the same symmetries as $H_X = \otimes_i X_i$ and its ground states—namely, the symmetries generated by $U_X = \otimes_i X_i$ and $U_Z = \otimes_i Z_i$, where the tensor products are taken over all qubits. As such, the Ginzburg-Landau theory cannot distinguish these two phases; indeed, because the ground states of H_{TC} appear disordered or liquid-like at short length-scales, they cannot be identified or characterized by any local order parameters. Instead, to capture the distinctive long-range correlations in such states, one must measure operators or quantities supported on an extensive number of qubits. These include, for example, the Wilson loop observables supported on large closed loops[75,171,169,69] and nonlocal topological invariants such as topological entanglement entropy[91,102].

2.3 Symmetry-Protected Topological Phases

The above framework, which defines a quantum phase of matter as an equivalence class of gapped Hamiltonians which can be adiabatically deformed into one another, allows us to study a wide range of quantum phases beyond the reach of Landau theory. However, it is still incomplete: indeed, with this definition, two Hamiltonians can belong to the same phase even if they have different symmetries—that is, even if they belong to different phases according to Landau's theory. To ac-

count for the effects of symmetries, it is thus necessary to adjust our equivalence relation.

As a first step, we require all Hamiltonians H belonging to the same phase to possess the same symmetries. Then, we can further refine our equivalence relation: in particular, we consider two Hamiltonians to belong to the same phase if and only if they can be adiabatically deformed into one another using transformations which commute with the global symmetry [40]. With this definition, each equivalent class becomes smaller, and there are more classes or quantum phases.

Several examples of such classes can be considered. First, with the new definition, short-range-entangled (trivial) states can belong to different equivalence classes if they have different broken symmetries—these correspond precisely to Landau's symmetry-breaking states [40]. Second, we find certain states which can be connected to the trivial states with generic local transformations, but not with symmetry-preserving transformations; we say these states possess "symmetry-protected topological" (SPT) order. Third, topologically-ordered states which do not break any symmetry may in general belong to distinct phases under the new equivalence relation. In what follows, we consider the second case, namely the SPT phases. In particular, we illustrate the intriguing physical properties which distinguish SPT phases from the trivial phase by examining a simple one-dimensional example with $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry.

1D *ZXZ* Model

Let us consider a family of Hamiltonians acting on a one-dimensional spin-1/2 chain with open boundary conditions:

$$H = H_{ZXZ} + hH_X \tag{2.7}$$

where

$$H_{ZXZ} = -J\sum_{i=1}^{N-2} Z_i X_{i+1} Z_{i+2}, \qquad H_X = -J\sum_{i=1}^{N} X_i. \tag{2.8}$$

For simplicity, we assume that the number N of qubits is even. This family of Hamiltonians possesses a global $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, which is generated by the operators $P_{\text{even(odd)}} = \prod_{i \text{ even(odd)}} X_i$. When $h \gg 1$, the ground state is a quantum paramagnet, with all spins aligned in the X-direction. In this case, the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry operators act trivially on the ground state.

Meanwhile, when $h = 0$, all terms in the Hamiltonian are mutually commuting, so that any ground state has $+1$ expectation value for all stabilizer operators $Z_i X_{i+1} Z_{i+2} = 1$. Upon concatenating a list of these stabilizers, one finds that the *string operators* $Z_1 X_2 X_4 \ldots X_{N-2} Z_{N-1}$ and $Z_2 X_3 X_5 \ldots X_{N-1} Z_N$ are both equal to $+1$. In this case, the global symmetries act locally on the left and right boundaries:

$$P_{\text{even}} = P^L_{\text{even}} P^R_{\text{even}} \tag{2.9}$$

$$P_{\text{odd}} = P^L_{\text{odd}} P^R_{\text{odd}} \tag{2.10}$$

where $P^L_{\text{even}} = Z_1$, $P^R_{\text{even}} = Z_{N-1} X_N$, $P^L_{\text{odd}} = X_1 Z_2$, and $P^R_{\text{odd}} = Z_N$ are fractionalized symmetry operators which commute with all stabilizers in the Hamiltonian. Moreover, because P^L_{even} (respectively, P^R_{even}) anticommutes with P^L_{odd} (P^R_{odd}), the ground state space of H_{ZXZ} is four-fold degenerate due to the existence of *edge modes*[161].

This distinction between the ground state degeneracies suggest that the $h = 0$ case and $h \gg 1$ case belong to two different quantum phases. Indeed, this Hamiltonian is exactly solvable via Jordan-Wigner transformation[135], which confirms the presence of a quantum phase transition at the point $h = 1$ from a SPT phase to the trivial paramagnetic phase. Mathematically, the two phases are distinguished by the *projective representation* of the symmetry group on the edges of the 1D chain, which is nontrivial in the SPT phase, and trivial in the paramagnetic phase[41]. Moreover, the nonlocal Jordan-Wigner transformation shows that, like the topological phases discussed in the preceding section, the SPT phase can only be characterized by nonlocal order parameters such as the

aforementioned string operators.

2.4 Challenge: Quantum Phase Recognition

In this chapter, we have discussed two particularly intriguing classes of exotic quantum matter, namely the topologically-ordered and SPT-ordered phases. In particular, we have found that these phases are both outside the reach of Landau's classification of quantum phases, and cannot be identified using local order parameters. Instead, they are characterized through the measurement of nonlocal quantities such as topological invariants or string and loop operators. However, the extensive support of such operators also causes them to decay rapidly when local fluctuations, such as coherent perturbations or incoherent noise processes, are present in the system. This makes the identification of such phases particularly difficult in experiments, where such fluctuations are commonly present.

These considerations motivate us to pose the problem of *quantum phase recognition* (QPR): can we efficiently determine whether a given input quantum state ρ_{in} belongs to a particular quantum phase of matter? In Chapters 4 and 5 of this thesis, we develop methods which utilize near-term quantum computers to solve QPR, which are substantially more efficient and accurate than existing approaches.

If you don't make mistakes, you're not working on hard enough problems. And that's a mistake.

Frank Wilczek

3

Quantum Error Correction

As discussed in Chapter 1, quantum computation has the potential to improve and accelerate a broad range of applications ranging from factoring and cryptography, to chemistry and drug discovery. However, these computational advantages are only seen when a large number of operations are needed—otherwise, the problems can likely be solved using modern-day supercomputers. In the quantum setting, this implies that many qubits and quantum gates are needed; for example, mil-

lions of logical gates are required to use Shor's algorithm to factor 2000-bit numbers (an estimated size which would have cryptographic implications, and cannot be performed classically)[57,58,159,177]. Meanwhile, current state-of-the-art experiments exhibit two-qubit gate fidelities ranging from 10^{-2} to 10^{-4}, due to decoherence effects and hardware imperfections[3,54,105,181]. As such, it is necessary to adopt quantum error correction techniques, which use redundancy to protect quantum information in the presence of errors or noise.

In a typical experiment, the quantum errors can be divided into two classes. The first class of errors consist of *memory errors*, where a qubit may undergo a process such as spontaneous emission and decay from one state to another even if it is not involved in any quantum gate; we will begin by considering only this class of errors. On the other hand, *operation errors* such as over- or under-rotation of laser pulses constitute a second class of quantum errors. This larger class of errors requires the implementation of *fault-tolerant quantum computation*, and will be discussed later in Section 3.3. Finally, we conclude by discussing resource estimates for implementing fault-tolerant quantum algorithms, and pose a problem of how to efficiently perform fault-tolerant quantum computation.

3.1 Classical Error Correction

Before addressing quantum memory errors, let us first review the basic concepts of classical error correction. In this case, we would like to store a single classical bit $\in \{0, 1\}$ for some time τ, during which a bit flip error that interchanges $0 \leftrightarrow 1$ may occur with probability p_e. The idea of classical error correction is then to use redundant coding to detect and correct any possible error. For example, our single bit (which we call the *logical bit*) can be encoded into three *physical bits* as follows:

$$0_L \equiv 000, \qquad 1_L \equiv 111. \tag{3.1}$$

Let us now consider, for example, the logical state 0_L after time τ. With probability $(1-p_e)^3$, the three physical bits will still be in the state 000. On the other hand, one of the bits may be flipped to 1, and the probability of being in of of these configurations is $3p_e(1-p_e)^2$. Similarly, the probability of having two of the three bits flipped is $3p_e^2(1-p_e)$, while the probability for all three bits undergoing the bit-flip error is p_e^3.

If the error probability p_e is small, classical error correction may be done after the storage time τ by measuring the three bits; if the same result is obtained for all three measurements, we know what the logical state is from Equation (3.1), while if we obtain different results, we use majority voting to decode the logical bit (e.g. $001 \mapsto 0_L, 101 \mapsto 1_L$). In this way, error correction can allow us to obtain the correct result with probability

$$p_{\text{correct}} = (1-p_e)^3 + 3p_e(1-p_e)^2 = 1 - 3p_e^2 + 2p_e^3. \tag{3.2}$$

Equivalently, the error rate for the logical bit (i.e., the *logical error rate*) is $1 - p_{\text{correct}} = 3p_e^2 - 2p_e^3$. Thus, the error correction allows us to reduce the error rate if $3p_e^2 - 2p_e^3 < p_e$, which occurs if $p_e < 1/2$.

Typically, the error rate scales linearly with the storage time τ: $p_e = c\tau$ for sufficiently small τ. To store a logical bit for a longer time t without incurring too large of an error rate, one can divide the total storage time t into N shorter time segments of duration $\tau = t/N$, and perform the above error correction procedure after each time segment. For large N, the probability of obtaining the correct result after time t will be

$$p_{\text{correct}}(t) \cong \left(1 - 3\left(\frac{ct}{N}\right)^2\right)^N. \tag{3.3}$$

Notice that $p_{\text{correct}}(t)$ can be made arbitrarily close to 1 by choosing a large enough value of N.

3.2 Ideas Behind Quantum Error Correction

For quantum error correction, we would like to use the same idea of redundant coding as in classical error correction. However, there are a few potential challenges:

1. Errors can be continuous (e.g. $|0\rangle \mapsto |0\rangle + \varepsilon|1\rangle$, $|1\rangle \mapsto |1\rangle - \varepsilon|0\rangle$ instead of just discrete bit-flips $|0\rangle \leftrightarrow |1\rangle$).
2. Due to the quantum no-cloning theorem [122], it is not possible to replicate an arbitrary quantum state multiple times.
3. There are multiple types of errors (e.g. bit-flip $|0\rangle \leftrightarrow |1\rangle$ or phase-flip $|0\rangle + |1\rangle \leftrightarrow |0\rangle - |1\rangle$).
4. Error detection needs to be performed without measuring all qubits and destroying the quantum superpositions.

To address the first challenge, we note that an arbitrary single-qubit error can be expressed as a linear combination of Pauli errors

$$U = 1 + \alpha X + \beta Y + \gamma Z, \tag{3.4}$$

so it suffices to consider the bit-flip and phase-flip errors X and Z (note $Y \propto XZ$ errors can be corrected if both X and Z errors are corrected). To address the final challenge, we will first consider the storage of a state

$$|\psi\rangle = c_0|0\rangle + c_1|1\rangle \tag{3.5}$$

with probability p_e of bit-flip errors $X|\psi\rangle$; we will later consider extending our framework to also address phase-flip errors.

To store $|\psi\rangle$ while protecting against bit-flip errors, we use three physical qubits in analogy to the classical repetition code discussed above, and introduce the logical encoding

$$|0_L\rangle = |000\rangle, \qquad |1_L\rangle = |111\rangle. \tag{3.6}$$

The subspace of Hilbert space spanned by the logical states $\{|0_L\rangle, |1_L\rangle\}$ is known as the *code subspace* and contains our logically encoded state $|\psi_L\rangle = c_0|0_L\rangle + c_1|1_L\rangle$. Notice that in general, $|\psi_L\rangle \neq |\psi\rangle^{\otimes 3}$, so this error-correction scheme does not violate the no-cloning theorem as discussed in the second challenge above.

After a storage time τ, the three-qubit state will, with probability $(1-p_e)^3$, remain as $|\psi_L\rangle$ (i.e. if no errors occur). On the other hand, the probability of a bit-flip error occurring on exactly one of the qubits $i \in \{1,2,3\}$ and resulting in the state $X_i|\psi_L\rangle$ is $3p_e(1-p_e)^2$. Likewise, the probability of two bit-flip errors occurring and producing the state $X_iX_j|\psi_L\rangle$ is $3p_e^2(1-p_e)$, while the probability of all three qubits undergoing bit-flip errors is p_e^3.

To detect and correct for these errors without collapsing the quantum state, we can perform collective measurements onto the code subspace using a projector

$$P_0 = |0_L\rangle\langle 0_L| + |1_L\rangle\langle 1_L|, \tag{3.7}$$

If we measure $P_0 = 1$, to leading order in the error probability, the state is the same as before; otherwise, a bit-flip error must have occurred. In the latter case, we measure the operators

$$P_i = X_i P_0 X_i \tag{3.8}$$

for $i = 1, 2, 3$. Because

$$P_i X_j |\psi_L\rangle = \delta_{ij} X_i |\psi_L\rangle, \tag{3.9}$$

the measurement outcomes for these operators can be used to determine which bit was flipped, and the appropriate error correction operator X_i can then be applied. Similar to the classical case, this approach works perfectly if only one qubit was flipped, but does not produce the correct result if two or three qubits underwent errors. Thus, after the error-correction procedure, the correct state $|\psi_L\rangle$ is obtained with probability

$$p_{\text{correct}} = 1 - 3p_e^2 + 2p_e^3, \tag{3.10}$$

and the error rate is reduced compared to the physical qubit error rate if $p_e < 1/2$. As discussed in the context of classical error correction, long storage times can be divided into short intervals and the quantum error-correction procedure can be applied between consecutive intervals to reduce the final error rate.

Notice that this approach naturally addresses the challenge of continuous errors: for example, if a coherent error

$$|\psi_L\rangle \mapsto |\psi_L\rangle + \eta X_i|\psi_L\rangle \tag{3.11}$$

occurred, measurement of P_0 and all P_i will force the system into $|\psi_L\rangle$ if $P_0 = 1$ or $X_i|\psi_L\rangle$ if $P_i = 1$, and the same correction steps can be applied as before. The physical understanding here is that the measurement of P_0 and P_i projects the state into one subspace of the three-qubit space—either the logical (code) subspace $\mathcal{H}_L$ if $P_0 = 1$, or onto one of the orthogonal subspaces if $P_i = 1$.

While the P_0 and P_i are theoretically the most straightforward to use for error detection in this setup, one can instead measure the simpler operators Z_1Z_2 and Z_2Z_3. This is because the operator Z_iZ_j compares the i^{th} and j^{th} qubits and produces $+1$ if the qubits are the same in the Z basis, and -1 if they are different. Thus, each of the four possible outcomes for these measurements will project the quantum state into one of the four subspaces as discussed above, and a corresponding error-correction operation can be applied if necessary.

3.2.1 Phase Errors

The logical qubit encoding of Equation (3.6) can be used to correct bit-flip (Pauli-X) errors, but it does not protect against phase-flip errors

$$|\psi\rangle = Z|\psi\rangle = c_0|0\rangle + c_1|1\rangle \mapsto c_0|0\rangle - c_1|1\rangle. \tag{3.12}$$

Because Z acts on the states $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ as $Z|\pm\rangle = |\mp\rangle$, a three-qubit code which protects against such Pauli-Z errors can be designed by using the above code in a rotated basis:

$$|0_L\rangle = |+++\rangle, \qquad |1_L\rangle = |---\rangle. \tag{3.13}$$

For this code, the error detection and correction steps work as before, except by measuring the operators X_1X_2 and X_2X_3 in place of Z_1Z_2 and Z_2Z_3.

While the code (3.13) now corrects for Z errors, it no longer corrects for the bit-flip errors. To address both kinds of errors together, we must introduce more qubits as discussed below.

3.2.2 Shor's Nine-Qubit Code

Peter Shor's trick to addressing both X and Z errors was to combine the protection from the three-qubit codes (3.6) and (3.13) into a nine-qubit code [146]

$$|+_L\rangle = \frac{1}{2^{3/2}}\left(|000\rangle + |111\rangle\right)^{\otimes 3}, \quad |-_L\rangle = \frac{1}{2^{3/2}}\left(|000\rangle - |111\rangle\right)^{\otimes 3}. \tag{3.14}$$

To correct bit-flip errors, one can measure the operators Z_1Z_2, Z_2Z_3, Z_4Z_5, Z_5Z_6, Z_7Z_8, and Z_8Z_9. Then, for example, the measurement outcome $Z_1Z_2 = -1, Z_2Z_3 = +1$ would signify that the first qubit underwent a bit-flip error. On the other hand, phase-flip errors can also be detected by mea-

suring the operators $X_1X_2X_3X_4X_5X_6$ and $X_4X_5X_6X_7X_8X_9$. For instance, the outcome $X_1X_2X_3X_4X_5X_6 = -1, X_4X_5X_6X_7X_8X_9 = +1$ would indicate that one of the first three qubits experienced a phase-flip error. Notice that this detection procedure cannot distinguish which of these three qubits had the error. However, this is not a problem: indeed, assuming only a single phase-flip error occurred, applying Z_i for any $i = 1$, 2, or 3 will result in the same (correct) logical state. Thus, Shor's nine-qubit code works to correct X errors by comparing the qubit values ($|0\rangle$ or $|1\rangle$) within blocks of three qubits using the Z_iZ_j measurements, and comparing the relative phases of adjacent blocks by measuring the products of Pauli-X operators. In this way, Shor's code can be used to correct arbitrary single-qubit errors.

To determine when Shor's code results in a smaller logical error rate compared to the original error rate, we assume that the physical qubits each undergo bit-flip and phase-flip errors with probability p. Then, the error probability for the resulting logical qubit is given by the sum of the probability for two bit-flip errors occurring within one block of three qubits ($9p^2$), and the probability for two of the three blocks to contain a phase-flip error ($3p^2$). Thus, Shor's code improves the error rate for storing a quantum state if $12p^2 < p$.

In general, the threshold error rate below which error correction can improve the logical qubit fidelity depends on both the choice of quantum error-correction code and the underlying error model (for example, if bit-flip and phase-flip errors do not occur with equal probability). In Chapter 4, we utilize our novel, machine-learning-inspired quantum algorithm to optimize this threshold error rate by tailoring the choice of error-correction code based on the specific error model at hand. In doing so, we find that our resultant codes significantly outperform traditional codes such as the nine-qubit code, so that our strategy can be used to facilitate experimental realizations of high-fidelity logical qubits.

3.2.3 Stabilizer Codes

Shor's nine-qubit quantum error-correction code is one example of a *stabilizer code*. In particular, if we define *stabilizer operators* $S_1, S_2, \ldots S_8$ as

$$S_1 = Z_1Z_2 \qquad S_2 = Z_2Z_3 \qquad S_3 = Z_4Z_5 \qquad S_4 = Z_5Z_6 \tag{3.15}$$

$$S_5 = Z_7Z_8 \qquad S_6 = Z_8Z_9 \tag{3.16}$$

$$S_7 = X_1X_2X_3X_4X_5X_6 \qquad S_8 = X_4X_5X_6X_7X_8X_9 \tag{3.17}$$

then it follows from the above discussion that the logical subspace of the nine-qubit code is the simultaneous $+1$ eigenspace of all stabilizer operators. Notice that the stabilizer operators are all tensor products of Pauli operators which commute with one another; as such, they generate an abelian subgroup of the nine-qubit Pauli group $\mathcal{P}_9$, where the n-qubit Pauli group $\mathcal{P}_n$ is defined as the group of operators which are n-qubit tensor products of 1, X, Y, or Z with an overall phase of ± 1 or $\pm i$. Any single-qubit error X_i or Z_i then anticommutes with a subset of the stabilizer operators, so that we can identify such an error when stabilizer operators within this subset are measured to be -1. Moreover, the nontrivial logical operators acting on the code subspace correspond to unitary operators which commute with all stabilizers, but cannot be constructed from a product of stabilizer operators. In general, such logical gates may be a complicated nine-qubit unitary operator; however, for some specific gates, the logical operation is *transversal*, meaning it can be decomposed into a tensor product of the corresponding physical gates. For example, the logical operator X_L which flips $|0_L\rangle \leftrightarrow |1_L\rangle$ in the nine-qubit code is given by

$$X_L = X_1X_2 \ldots X_9, \tag{3.18}$$

while the logical Z operator is given by

$$Z_L = Z_1 Z_2 ... Z_9. \tag{3.19}$$

More generally, we can consider stabilizer operators S_i which belong to an abelian subgroup of the n-qubit Pauli group, $S \subsetneq \mathcal{P}_n$. Because $\left[S_i, S_j\right] = 0$ for all $S_i, S_j \in S$, the operators $S_i \in S$ are simultaneously diagonalizable. We then say that a state $|\psi\rangle$ is *stabilized* by S if for all stabilizer operators $S_i \in S$, $S_i|\psi\rangle = |\psi\rangle$[64]. In general, if the subgroup S has l generators for n physical qubits, each subspace labeled by a set of eigenvalues $s_1, ..., s_{2^l}$ of $S_1, ..., S_{2^l}$ has dimensionality 2^{n-l}. The logical subspace (i.e., the space with all $s_i = +1$) can therefore be used to encode $n - l$ logical qubits. Then, errors can likewise be detected through -1 measurements of stabilizer operators S_i, and logical operators are also given by unitary operators which commute with all $S_i \in S$.

This stabilizer-based framework can be used to identify and characterize many quantum error-correcting codes. In this thesis, we consider three important stabilizer codes. The first is known as the seven-qubit Steane code[150], which has six stabilizer operators

$$\begin{aligned} g_1 &= IIIXXXX & g_2 &= IXXIIXX & g_3 &= XIXIXIX \\ g_4 &= IIIZZZZ & g_5 &= IZZIIZZ & g_6 &= ZIZIZIZ, \end{aligned} \tag{3.20}$$

and whose code subspace is spanned by the logical states

$$\begin{aligned} |0_L\rangle = &\frac{1}{2\sqrt{2}} \left(|0000000\rangle + |1010101\rangle + |0110011\rangle + |1100110\rangle \right. \\ &\left. + |0001111\rangle + |1011010\rangle + |0111100\rangle + |1101001\rangle\right) \end{aligned} \tag{3.21}$$

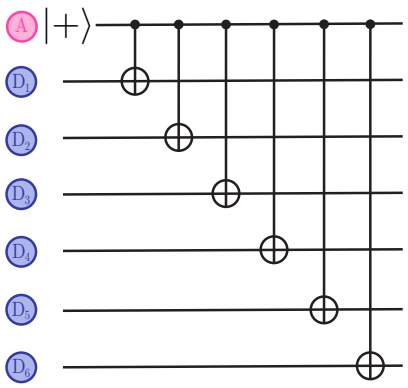


Figure 3.1: Circuit to measure the operator $X_1 \ldots X_6$ for the nine-qubit Shor code. The ancilla qubit (pink, A) is prepared in the superposition state $|\psi_A\rangle = |+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, and controlled-NOT gates (with A as the control qubit) are performed between A and each data qubit (blue, D_i) for $i = 1, 2, \ldots 6$. The measurement of A in the X basis then gives the desired eigenvalue.

$$\begin{aligned} |1\rangle_L = &\frac{1}{2\sqrt{2}} (|1111111\rangle + |0101010\rangle + |1001100\rangle + |0011001\rangle \\ &+ |1110000\rangle + |0100101\rangle + |1000011\rangle + |0010110\rangle) . \end{aligned} \tag{3.22}$$

Second, we consider the X-basis repetition code introduced in Section 3.2.1, where the stabilizer operators are X_1X_2 and X_2X_3. Finally, we consider the surface code[92,22], where the physical qubits are placed on a two-dimensional square lattice. In the surface code, the logical subspace is precisely the ground state space of the toric code Hamiltonian H_{TC} introduced in Section 2.2, and the stabilizer operators are given by the vertex and plaquette terms of Equation (2.6); this logical subspace is non-trivial when the lattice is placed on an underlying surface with non-trivial topology, such as a torus or an annulus. In Chapter 5, we discover that stabilizer-based ideas from the surface code can be used to facilitate certain solutions of the quantum phase recognition problem introduced in Section 2.4.

When using one of these stabilizer codes to implement quantum error correction, one of the most important tasks is to measure stabilizer operators, which are products of Pauli operators. For example, to detect and correct for Pauli-Z errors in the nine-qubit code, we must measure the stabi-

lizer operator $X_1 \ldots X_6$. Such an operator can be measured by preparing an ancilla qubit A in the state $|\psi_A\rangle = |+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and performing CNOT gates between the ancilla and each data qubit i, as shown in the circuit of Figure 3.1. Upon measuring the ancilla qubit in the X basis and projecting it into one of the states $|0\rangle \pm |1\rangle$, we obtain the desired eigenvalue $X_1 \ldots X_6 = \pm 1$. Similar circuits can be used to measure the other operators for quantum error correction such as $X_4 \ldots X_9$ and $Z_1 Z_2$, $Z_2 Z_3$, etc.; in total, we find that 24 CNOT gates are needed to measure all operators for one round of quantum error correction.

3.3 Challenge: Fault-Tolerant Quantum Computation

As mentioned in the beginning of this chapter, in addition to memory errors which can occur on any qubit, we must also consider operational errors which occur on qubits involved in quantum gates. These operational errors arise mainly from two sources: First, when multi-qubit entangling gates are involved, an error on any one of the qubits may be propagated to all qubits involved in the gate. For example, if a Pauli-X error is present on the control qubit of a CNOT gate, it will spread to Pauli-X errors on both the control and target qubits after the CNOT:

$$U_{\text{CNOT}} X_1 |\psi\rangle = (U_{\text{CNOT}} X_1 U_{\text{CNOT}}^\dagger) U_{\text{CNOT}} |\psi\rangle = X_1 X_2 U_{\text{CNOT}} |\psi\rangle. \tag{3.23}$$

A second type of operational error arises when the gate itself is imperfect, for instance if an entangling gate U is implemented instead as some quantum channel $\mathcal{E}$. For a two-qubit gate, such an error can also be expressed as a linear combination of two-qubit Pauli operators.

Unfortunately, errors such as (3.23) can be particularly destructive—for example, if a circuit creates two errors within a single logical qubit in the nine-qubit code, then the original state can no longer be recovered. However, if only single-qubit errors are produced, these operational errors can be addressed in the same way as memory errors. These observations lead to the concept of *fault-*

tolerant quantum computation, where the idea is to create conditions such that at most one error per logical qubit is introduced at each computational step, and any such error gets corrected at the subsequent step. That is, given the probability p of failure of each individual component (e.g. a CNOT gate), the evolution must be constructed carefully such that the probability to introduce two errors within the same logical qubit is Cp^2 for some constant factor C.

Due to the possibility of propagation of errors following Equation (3.23) during error-detection circuits such as Figure 3.1, enforcing this property can often require a prohibitively large overhead in terms of the number of additional qubits required, or the number of gates needed within each logical component. In Chapter 6 of this thesis, we propose methods to address this important challenge for building scalable, error-corrected quantum computers. In particular, by introducing *hardware-efficient* approaches which utilize the features of the underlying hardware platform to address its possible errors, we substantially reduce the resource cost of fault-tolerant quantum computation through protocols which can be implemented in current or near-term neutral-atom-based quantum computers.

Learning is not attained by chance, it must be sought for with ardor and attended to with diligence.

Abigail Adams

4

Quantum Convolutional Neural Networks

Neural network-based machine learning has recently proven successful for many complex applications ranging from image recognition to precision medicine. As such, given the intricate complexity of quantum many-body physics, it is natural to consider utilizing such methods for solving problems such as the quantum phase recognition problem posed in Chapter 2. However, the direct application of machine learning methods to quantum physics is challenging due to the ex-

ponential complexity of many-body systems. Motivated by recent advances in realizing quantum information processors, we introduce and analyze a quantum circuit-based algorithm inspired by convolutional neural networks, a highly effective model in machine learning. Our quantum convolutional neural network (QCNN) uses only $O(\log(N))$ variational parameters for input sizes of N qubits, allowing for its efficient training and implementation on realistic, near-term quantum devices. To explicitly illustrate its power for solving quantum phase recognition, we show that QCNN can accurately recognize quantum states associated with a 1D symmetry-protected topological (SPT) phase, with performance surpassing existing approaches. We further demonstrate that QCNN can be used to devise a quantum error correction scheme optimized for a given, unknown error model that significantly outperforms known quantum codes of comparable complexity. Finally, we discuss potential experimental realizations of QCNN in near-term quantum devices.

4.1 Introduction and Motivations

The complex nature of quantum many-body systems motivates using machine learning techniques to analyze them. Indeed, large-scale neural networks have successfully solved classically difficult problems such as image recognition or improving classical error correction [99], and their architectures have been related to various physical concepts [106,115]. As such, a number of recent works have used neural networks for studying properties of quantum many-body systems [30,157,32,167,105,182,112]. However, the direct application of these classical algorithms is challenging for intrinsically quantum problems, which take quantum states or processes as inputs. This is because the extremely large many-body Hilbert space hinders efficiently translating such problems into a classical framework without performing exponentially difficult quantum state or process tomography [70,100].

Recent experimental progress towards realizing quantum information processors [95,118,49,5] has led to proposals for using quantum computers to enhance conventional machine learning tasks [15,52,55,81].

Motivated by such developments, we introduce and analyze a machine learning-inspired quantum circuit model—the quantum convolutional neural network (QCNN)—and demonstrate its ability to solve important classes of intrinsically quantum many-body problems. The first class of problems we consider is *quantum phase recognition* (introduced in Chapter 2), which asks whether a given input quantum state ρ_{in} belongs to a particular quantum phase of matter. In contrast to many existing schemes based on tensor network descriptions [80,149,90], we assume ρ_{in} is prepared in a physical system without direct access to its classical description. The second class, *quantum error correction (QEC) optimization* (introduced in Chapter 3), asks for an optimal QEC code for a given, *a priori* unknown error model such as dephasing or potentially correlated depolarization in realistic experimental settings. We provide both theoretical insight and numerical demonstrations for the successful application of QCNN to these important problems, and show its feasibility for near-term experimental implementation.

4.2 QCNN Circuit Model

Convolutional neural networks (CNNs) provide a successful machine learning architecture for classification tasks such as image recognition [98,93,99]. A CNN generally consists of a sequence of different (interleaved) layers of image processing; in each layer, an intermediate 2D array of pixels, called a feature map, is produced from the previous one (Figure 4.1a). (More generally, CNN layers connect *volumes* of multiple feature maps to subsequent volumes; for simplicity, we consider only a single feature map per volume and leave the generalization to future works.) The convolution layers compute new pixel values $x_{ij}^{(\ell)}$ from a linear combination of nearby ones in the preceding map $x_{i,j}^{(\ell)} = \sum_{a,b=1}^{w} w_{a,b} x_{i+a,j+b}^{(\ell-1)}$, where the weights $w_{a,b}$ form a $w \times w$ matrix. Pooling layers reduce feature map size, e.g. by taking the maximum value from a few contiguous pixels, and are often followed by application of a nonlinear (activation) function. Once the feature map size becomes

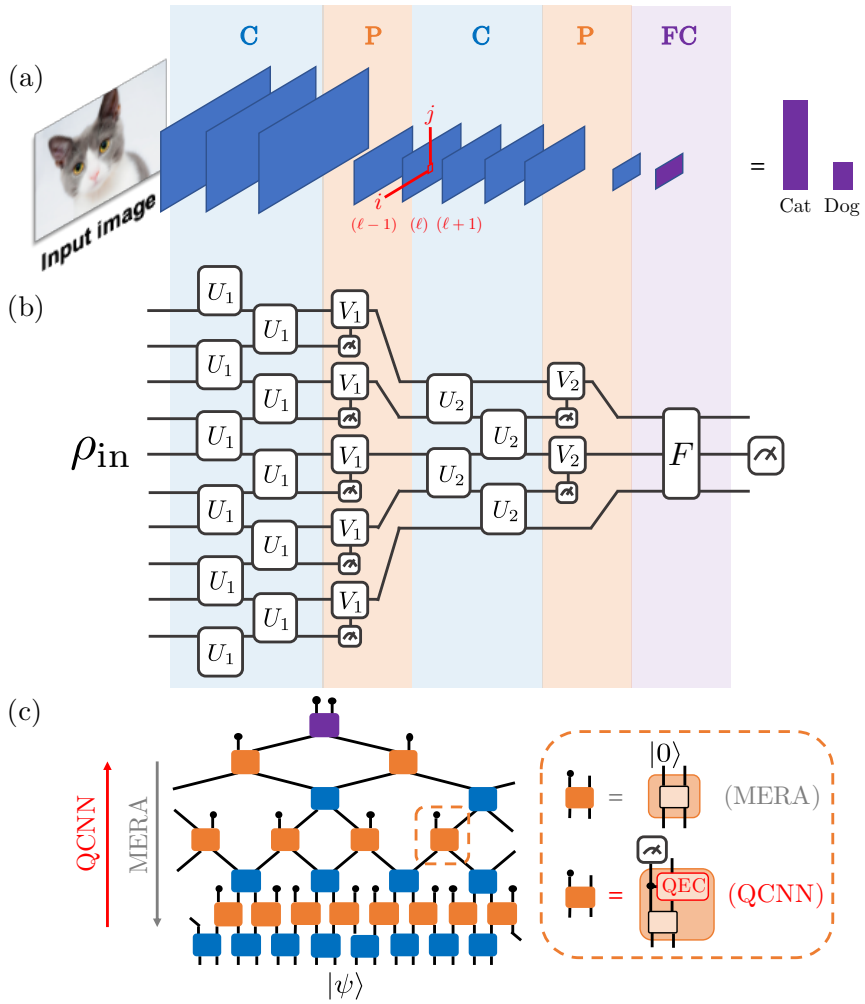


Figure 4.1: The concept of QCNN. (a) Simplified illustration of classical CNNs. A sequence of image processing layers—convolution (C), pooling (P), and fully connected (FC)—transforms an input image into a series of feature maps (blue rectangles), and finally into an output probability distribution (purple bars). (b) QCNNs inherit a similar layered structure. (c) QCNN and MERA share the same circuit structure, but run in reverse directions.

sufficiently small, the final output is computed from a function that depends on all remaining pixels (fully connected layer). The weights and fully connected function are optimized by training on large datasets. In contrast, variables such as the number of convolution and pooling layers and the size w of the weight matrices (known as hyperparameters) are fixed for a specific CNN[99]. CNN's key properties are thus translationally invariant convolution and pooling layers, each characterized by a constant number of parameters (independent of system size), and sequential data size reduction (i.e., a hierarchical structure).

Motivated by this architecture, we introduce a quantum circuit model (QCNN) extending these key properties to the quantum domain (Figure 4.1b). The circuit's input is an unknown quantum state ρ_{in}. A convolution layer applies a single quasi-local unitary (U_i) in a translationally-invariant manner for finite depth. For pooling, a fraction of qubits are measured, and their outcomes determine unitary rotations (V_j) applied to nearby qubits. Hence, nonlinearities in QCNN arise from reducing the number of degrees of freedom. Convolution and pooling layers are performed until the system size is sufficiently small; then, a fully connected layer is applied as a unitary F on the remaining qubits. Finally, the outcome of the circuit is obtained by measuring a fixed number of output qubits. As in the classical case, QCNN hyperparameters such as the number of convolution and pooling layers are fixed, and the unitaries themselves are learned.

A QCNN to classify N-qubit input states is thus characterized by $O(\log(N))$ parameters. This corresponds to doubly exponential reduction compared to a generic quantum circuit-based classifier[55] and allows for efficient learning and implementation. For example, given classified training data $\{(|\psi_\alpha\rangle, y_\alpha) : \alpha = 1, \ldots, M\}$, where $|\psi_\alpha\rangle$ are input states and $y_\alpha = 0$ or 1 are corresponding binary classification outputs, one could compute the mean-squared error

$$\text{MSE} = \frac{1}{2M}\sum_{\alpha=1}^{M}(y_i - f_{\{U_i,V_j,F\}}(|\psi_\alpha\rangle))^2. \tag{4.1}$$

Here, $f_{\{U_i,V_j,F\}}(|\psi_\alpha\rangle)$ denotes the expected QCNN output value for input $|\psi_\alpha\rangle$. Learning then consists of initializing all unitaries and successively optimizing them until convergence, e.g. via gradient descent.

To gain physical insight into the mechanism underlying QCNNs and motivate their application to the problems under consideration, we now relate our circuit model to two well-known concepts in quantum information theory—the multiscale entanglement renormalization ansatz [165] (MERA) and quantum error correction (QEC). The MERA framework provides an efficient tensor network representation of many classes of interesting many-body wavefunctions, including those associated with critical systems [165,1,126]. A MERA can be understood as a quantum state generated by a sequence of unitary and isometry layers applied to an input state (e.g. $|00\rangle$). While both types of layers apply quasilocal unitary gates, each isometry layer first introduces a set of new qubits in a predetermined state, e.g. $|0\rangle$ (Figure 4.1c). This exponentially growing, hierarchical structure allows for the long-range correlations associated with critical systems. The QCNN circuit has similar structure, but runs in the reverse direction. Hence, for any given state $|\psi\rangle$ with a MERA representation, there is always a QCNN that recognizes $|\psi\rangle$ with deterministic measurement outcomes; one such QCNN is simply the inverse of the MERA circuit.

For input states other than $|\psi\rangle$, however, such a QCNN does not generally produce deterministic measurement outcomes. These additional degrees of freedom distinguish QCNN from MERA. Specifically, we can identify the measurements as syndrome measurements in QEC [129], which determine error correction unitaries V_j to apply to the remaining qubit(s). Thus, a QCNN circuit with multiple pooling layers can be viewed as a combination of MERA — an important variational ansatz for many-body wavefunctions — and nested QEC — a mechanism to detect and correct local quantum errors without collapsing the wavefunction. This makes QCNN a powerful architecture to classify input quantum states or devise novel QEC codes. In particular, for QPR, the QCNN can provide a MERA realization of a representative state $|\psi_0\rangle$ in the target phase. Other input states

within the same phase can be viewed as $|\psi_0\rangle$ with local errors, which are repeatedly corrected by the QCNN in multiple layers. In this sense, the QCNN circuit can mimic renormalization-group (RG) flow, a methodology which successfully classifies many families of quantum phases [135].

4.3 Detecting a 1D Symmetry-Protected Topological Phase

We first demonstrate the potential of QCNN by applying it to QPR in a class of one-dimensional many-body systems. Specifically, we consider a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry-protected topological (SPT) phase $\mathcal{P}$, a phase containing the $S = 1$ Haldane chain [73]. We study ground states $\{|\psi_G\rangle\}$ of a family of Hamiltonians on a spin-1/2 chain with open boundary conditions (see Section 2.3):

$$H = -J \sum_{i=1}^{N-2} Z_i X_{i+1} Z_{i+2} - h_1 \sum_{i=1}^{N} X_i - h_2 \sum_{i=1}^{N-1} X_i X_{i+1}. \tag{4.2}$$

X_i, Z_i are Pauli operators for the spin at site i, and the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry is generated by $P_{\text{even(odd)}} = \prod_{i \in \text{even(odd)}} X_i$. Figure 4.2a shows the phase diagram as a function of $(h_1/J, h_2/J)$. When $h_2 = 0$, the Hamiltonian is exactly solvable via Jordan-Wigner transformation [135], confirming that $\mathcal{P}$ is characterized by nonlocal order parameters. When $h_1 = h_2 = 0$, all terms are mutually commuting, and a ground state is the 1D cluster state. Our goal is to identify whether an given, unknown ground state drawn from the phase diagram belongs to $\mathcal{P}$.

As an example, we first present an exact, analytical QCNN circuit that recognizes $\mathcal{P}$ (Figure 4.2b). The convolution layers involve controlled-phase gates as well as Toffoli gates with controls in the X-basis, and pooling layers perform phase-flips on remaining qubits when one adjacent measurement yields $X = -1$. This convolution-pooling unit is repeated d times, where d is the QCNN depth. The fully connected layer measures $Z_{i-1} X_i Z_{i+1}$ on the remaining qubits. Figure 4.2c shows the QCNN output for a system of $N = 135$ spins and $d = 1, \ldots, 4$ along $h_2 = 0.5J$, obtained using matrix product state simulations. As d increases, the measurement outcomes show sharper

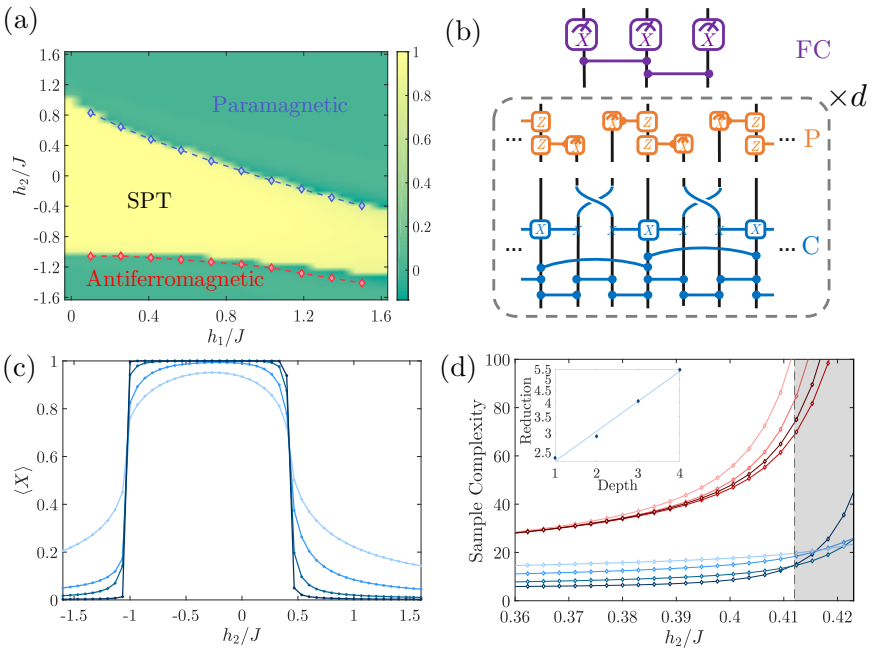


Figure 4.2: Application to quantum phase recognition. (a) The phase diagram of the Hamiltonian in the main text. The phase boundary points (blue and red diamonds) are extracted from infinite size DMRG numerical simulations, while the color represents the output from the exact QCNN circuit for input size $N = 45$ spins (see Appendix A.1). (b) Exact QCNN circuit to recognize a $\mathbb{Z}_2 \times \mathbb{Z}_2$ SPT phase. Blue line segments represent controlled-phase gates, blue three-qubit gates are Toffoli gates with the control qubits in the X basis, and orange two-qubit gates flip the target qubit's phase when the X measurement yields -1. The fully connected layer applies controlled-phase gates followed by an X_i projection, effectively measuring $Z_{i-1}X_iZ_{i+1}$. (c) Exact QCNN output along $h_1 = 0.5J$ for $N = 135$ spins, $d = 1, \ldots, 4$. (d) Sample complexity of QCNN at depths $d = 1, \ldots 4$ (blue) versus SOPs of length $N/2, N/3, N/5,$ and $N/6$ (red) to detect the SPT/paramagnet phase transition along $h_1 = 0.5J$ for $N = 135$ spins. The critical point is identified as $h_2/J = 0.423$ using infinite size DMRG (bold line). Darkening colors show higher QCNN depth or shorter string lengths. In the shaded area, the correlation length exceeds the system size and finite-size effects can considerably affect our results. Inset: The ratio of SOP sample complexity to QCNN sample complexity is plotted as a function of depth d on a logarithmic scale for $h_1/J = 0.3918$. In the numerically accessible regime, this reduction of sample complexity scales exponentially as $1.73e^{0.28d}$ (trendline).

changes around the critical point, and the output of a $d = 2$ circuit already reproduces the phase diagram with high accuracy (Figure 4.2a). This QCNN can also be used for other Hamiltonian models belonging to the same phase, such as the $S = 1$ Haldane chain[73], as discussed at the end of this section. More details on the construction of this exact, analytical QCNN circuit can be found in Appendix A.3.

4.3.1 Sample Complexity

The performance of a QPR solver can be quantified by sample complexity[70]: what is the expected number of copies of the input state required to identify its quantum phase? We demonstrate that the sample complexity of our exact QCNN circuit is significantly better than that of conventional methods. In principle, $\mathcal{P}$ can be detected by measuring a nonzero expectation value of string order parameters (SOP)[71,128] such as

$$\mathcal{S}_{ab} = Z_a X_{a+1} X_{a+3} \ldots X_{b-3} X_{b-1} Z_b \tag{4.3}$$

where $a < b$. In practice, however, the expectation values of SOP vanish near the phase boundary due to diverging correlation length [128]; since quantum projection noise is maximal in this vicinity, many experimental repetitions are required to affirm a nonzero expectation value. In contrast, the QCNN output is much sharper near the phase transition, so fewer repetitions are required.

Quantitatively, given some $|\psi_{\text{in}}\rangle$ and SOP S, a projective measurement of S can be modeled as a (generalized) Bernoulli random variable, where the outcome is 1 with probability $p = (\langle\psi_{\text{in}}| S |\psi_{\text{in}}\rangle + 1)/2$ and -1 with probability $1 - p$ (since $S^2 = 1$); after M binary measurements, we estimate p. $p > p_0 = 0.5$ signifies $|\psi_{\text{in}}\rangle \in \mathcal{P}$. We define the sample complexity M_{min} as the minimum M to test

whether $p > p_0$ with 95% confidence using an arcsine variance-stabilizing transformation [26]:

$$M_{\min} = \frac{1.96^2}{(\arcsin\sqrt{p} - \sqrt{\arcsin p_0})^2}. \tag{4.4}$$

Similarly, the sample complexity for a QCNN can be determined by replacing $\langle\psi_{\text{in}}| S |\psi_{\text{in}}\rangle$ by the QCNN output expectation value in the expression for p.

Figure 4.2d shows the sample complexity for the QCNN at various depths and SOPs of different lengths. Clearly, QCNN requires substantially fewer input copies throughout the parameter regime, especially near criticality. More importantly, although the SOP sample complexity scales independently of string length, the QCNN sample complexity consistently improves with increasing depth and is only limited by finite size effects in our simulations. In particular, compared to SOPs, QCNN reduces sample complexity by a factor which scales exponentially with the QCNN's depth in numerically accessible regimes (inset). Such scaling arises from the iterative QEC performed at each depth and is not expected from any measurements of simple (potentially nonlocal) observables. We show in Appendix A.2 that our QCNN circuit measures a *multiscale* string order parameter—a sum of products of exponentially many different SOPs which remains sharp up to the phase boundary.

4.3.2 MERA and QEC

Additional insights into the QCNN's performance are revealed by interpreting it in terms of MERA and QEC. In particular, our QCNN is specifically designed to contain the MERA representation of the 1D cluster state ($|\psi_0\rangle$) such that it becomes a stable fixed point. When $|\psi_0\rangle$ is fed as input, each convolution-pooling unit produces the same state $|\psi_0\rangle$ with reduced system size in the unmeasured qubits, while yielding deterministic outcomes ($X = 1$) in the measured qubits. The fully connected layer measures the SOP for $|\psi_0\rangle$. When an input wavefunction is perturbed away from $|\psi_0\rangle$, our

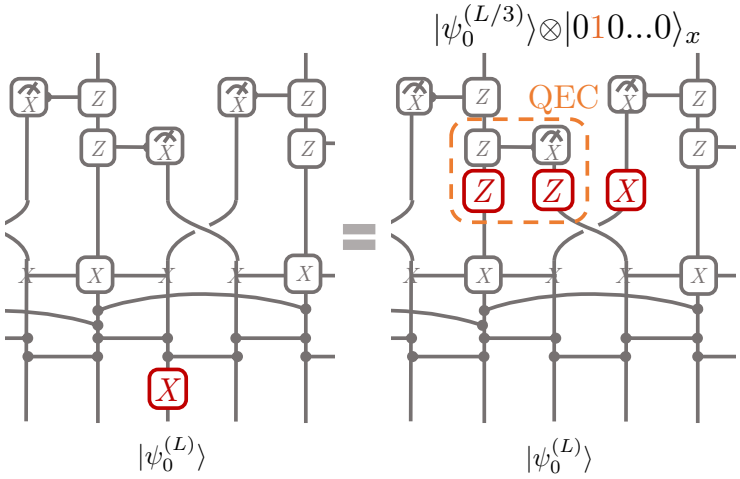


Figure 4.3: MERA and QEC in the QCNN circuit. The convolution-pooling unit of our circuit identifies and corrects single-qubit X errors acting on the cluster state, while reducing the system size by a factor of 3. This process resembles the combination of MERA and QEC.

QCNN corrects such "errors." For example, if a single X error occurs, the first pooling layer identifies its location, and controlled unitary operations correct the error propagated through the circuit (Fig. 4.3). Similarly, if an initial state has multiple, sufficiently separated errors (possibly in coherent superpositions), the error density after several iterations of convolution and pooling layers will be significantly smaller [178]. If the input state converges to the fixed point, our QCNN classifies it into the SPT phase with high fidelity. Clearly, this mechanism resembles the classification of quantum phases based on renormalization-group (RG) flow. This theoretical understanding also enables the construction of QCNN circuits for more generic QPR problems (see Appendix A.4 and Figure A.1).

4.3.3 Obtaining QCNN from Training Procedure

Having analytically illustrated the computational power of the QCNN circuit model, we now demonstrate how a QCNN for $\mathcal{P}$ can also be obtained using the learning procedure. Details of the QCNN's hyperparameters can be found in Appendix A.5 and Figure A.2. Initially, all unitaries

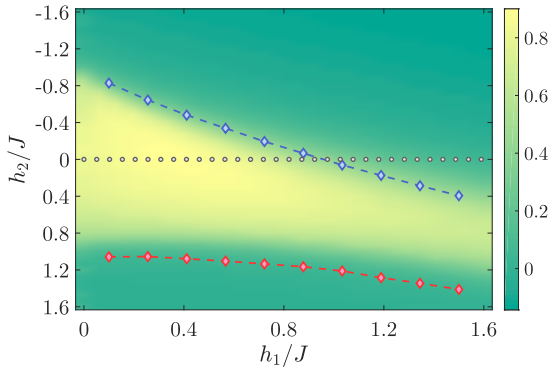


Figure 4.4: Output of a trained QCNN. We numerically optimize our QCNN for a system of $N = 15$ spins and depth $d = 1$ starting from random initial values. The training data points are 40 equally spaced points $h_1 \in [0, 2]$ along the line $h_2 = 0$ where the Hamiltonian is solvable by Jordan-Wigner transformation (e.g. gray dots). The blue and red diamonds are phase boundary points extracted from infinite size DMRG numerical simulations, while the colors represent the expected QCNN output value.

are set to random values. Because classically simulating our training procedure requires expensive computational resources, we focus on a relatively small system with $N = 15$ spins and QCNN depth $d = 1$; there are a total of 1309 parameters to be learned (see Appendix A.5). Our training data consists of 40 evenly spaced points along the line $h_2 = 0$, where the Hamiltonian is exactly solvable by Jordan-Wigner transformation. Using gradient descent with the mean-squared error function (4.1), we iteratively update the unitaries until convergence (see Appendix A.5). The classification output of the resulting QCNN for generic h_2 is shown in Figure 4.4. Remarkably, this QCNN accurately reproduces the 2D phase diagram over the entire parameter regime, even though the model was trained only on samples from a set of solvable points which does not even cross the lower phase boundary.

This example illustrates how the QCNN structure avoids overfitting to training data with its exponentially reduced number of parameters. While the training dataset for this particular QPR problem consists of solvable points, more generally, such a dataset can be obtained by using traditional methods (e.g. measuring SOPs) to classify representative states that can be efficiently generated either numerically or experimentally [142,61].

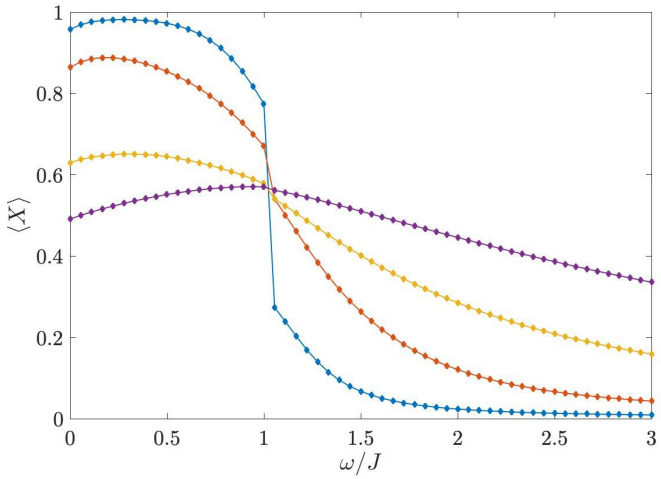


Figure 4.5: Exact QCNN output (at depth $d = 1, ...4$) for the Haldane chain Hamiltonian with $N = 54$ spins.

4.3.4 QCNN for the $S = 1$ Haldane Chain

As discussed at the beginning of this section, the (spin-1/2) 1D cluster state belongs to an SPT phase protected by $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, a phase which also contains the celebrated $S = 1$ Haldane chain[73]. It is thus natural to ask whether this circuit can be used to detect the phase transition between the Haldane phase and an $S = 1$ paramagnetic phase, which we numerically demonstrate here.

The one-parameter family of Hamiltonians we consider for the Haldane phase is defined on a one-dimensional chain of N spin-1 particles with open boundary conditions[73]:

$$H_{\text{Haldane}} = J \sum_{j=1}^{N} \mathbf{S_j} \cdot \mathbf{S_{j+1}} + \omega \sum_{j=1}^{N} (S_j^x)^2 \tag{4.5}$$

In this equation, $\mathbf{S_j}$ denotes the vector of $S = 1$ spin operators at site j. The system is protected by a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry generated by global π-rotations of every spin around the X and Y axes:

$R_x = \prod_j e^{i\pi S_j^x}$, $R_y = e^{i\pi S_j^y}$. When ω is zero or small compared to J, the ground state belongs to the SPT phase, but when ω/J is sufficiently large, the ground state becomes paramagnetic[73].

To apply our QCNN circuit to this Haldane phase, we must first identify a quasi-local isometric map U between the two models, because their representations of the symmetry group are distinct. More specifically, since the cluster model has a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry generated by $P_{\text{even(odd)}} = \prod_{i\in\text{even(odd)}} X_i$, we require $UR_xU^\dagger = P_{\text{odd}}$ and $UR_yU^\dagger = P_{\text{even}}$. Such a map can be constructed following Ref.[161]. Intuitively, it extends the local Hilbert space of a spin-1 particle by introducing a spin singlet state $|s\rangle$ and mapping it to a pair of spin-1/2 particles: $|x\rangle \mapsto |+-\rangle$, $|y\rangle \mapsto -|-+\rangle$, $|z\rangle \mapsto -i|--\rangle$, $|s\rangle \mapsto |++\rangle$. Here, $|\pm\rangle$ denote the ±1 eigenstates of the (spin-1/2) Pauli matrix X. $|\mu\rangle$ denotes a spin-1 state defined by $R_\nu|\mu\rangle = (-1)^{\delta_{\mu,\nu}+1}|\mu\rangle$ ($\mu,\nu \in \{x,y,z\}$). The QCNN circuit for the Haldane chain thus consists of applying U followed by the circuit presented in the main text.

Figure 4.5 shows the QCNN output for an input system of $N = 54$ spin-1 particles at depths $d = 1,\ldots,4$, obtained using matrix product state simulations with bond dimension $D = 160$. For this system size, we numerically identified the critical point as $\omega/J = 1.035 \pm 0.005$, by using DMRG to obtain the second derivative of energy density as a function of ω and J. The QCNN provides accurate identification of the phase transition.

4.4 Optimizing Quantum Error Correction

As seen in the previous example, the QCNN's architecture enables one to perform effective QEC. We next leverage this feature to design a new QEC code itself that is optimized for a given error model. More specifically, any QCNN circuit (and its inverse) can be viewed as a decoding (encoding) quantum channel between the physical input qubits and the logical output qubit. The encoding scheme introduces sets of new qubits in a predetermined state, e.g. $|0\rangle$, while the decoding scheme performs measurements (Fig. 5a). Given a error channel $\mathcal{N}$, our aim is therefore to maxi-

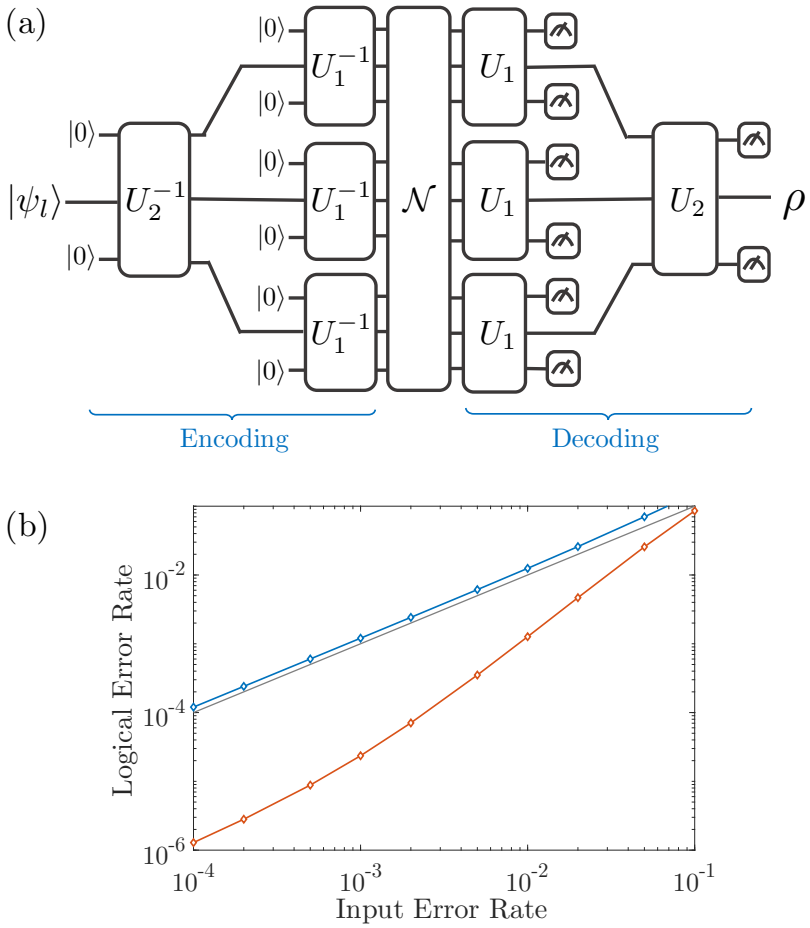


Figure 4.6: QCNN for optimizing quantum error correction. (a) Schematic diagram for using QCNNs to optimize QEC. The inverse QCNN encodes a single logical qubit $|\psi_l\rangle$ into 9 physical qubits, which undergo noise $\mathcal{N}$. QCNN then decodes these to obtain the logical state ρ. Our aim is to maximize $\langle\psi_l|\rho|\psi_l\rangle$. (b) Logical error rate of Shor code (blue) versus a learned QEC code (orange) in a correlated error model. The input error rate is defined as the sum of all probabilities p_μ and p_{xx}. The Shor code has worse performance than performing no error correction at all (identity, gray line), while the optimized code can still significantly reduce the error rate.

mize the recovery fidelity

$$f_q = \sum_{|\psi_l\rangle \in \{|\pm x, y, z\rangle\}} \langle \psi_l | \, \mathcal{M}_q^{-1}(\mathcal{N}(\mathcal{M}_q(|\psi_l\rangle \langle \psi_l|))) \, |\psi_l\rangle \, , \tag{4.6}$$

where $\mathcal{M}_q(\mathcal{M}_q^{-1})$ is the encoding (decoding) scheme generated by a QCNN circuit, and $|\pm x, y, z\rangle$ are the ± 1 eigenstates of the Pauli matrices. Thus, our method simultaneously optimizes both encoding and decoding schemes, while ensuring their efficient implementation in realistic systems. Importantly, the variational optimization can be carried out with a unknown $\mathcal{N}$ since f_q can be evaluated experimentally.

To illustrate the potential of this procedure, we consider a two-layer QCNN with $N = 9$ physical qubits and 126 variational parameters (Figure 4.6a and Appendix A.6). This particular architecture includes the nested (classical) repetition codes and the 9-qubit Shor code [146]; in the following, we compare our performance to the better of the two. We consider three different input error models: (1) independent single-qubit errors on all qubits with equal probabilities p_μ for $\mu = X$, Y, and Z errors or (2) anisotropic probabilities $p_x \neq p_y = p_z$, and (3) independent single-qubit anisotropic errors with additional two-qubit correlated errors $X_i X_{i+1}$ with probability p_{xx}.

Upon initializing all QCNN parameters to random values and numerically optimizing them to maximize f_q, we find that our model produces the same logical error rate as known codes in case (1), but can reduce the error rate by a constant factor of up to 50% in case (2), depending on the specific input error probability ratios (see Appendix A.6 and Figure A.3). More drastically, in case (3), the optimized QEC code performs significantly better than known codes (Figure 4.6b). Specifically, because the Shor code is only guaranteed to correct arbitrary single-qubit errors, it performs even worse than using no error correction, while the optimized QEC code performs much better. This example demonstrates the power of using QCNNs to obtain and optimize new QEC codes for realistic, a priori unknown error models.

4.5 Experimental Realizations

Our QCNN architecture can be efficiently implemented on several state-of-the-art experimental platforms. The key ingredients for realizing QCNNs include the efficient preparation of quantum many-body input states, the application of two-qubit gates at various length scales, and projective measurements. As in stabilizer-based QEC, the measurements of intermediate qubits and feed-forwarding can be replaced by controlled two-qubit unitary operations so that measurements are only performed at the end of an experimental sequence. These capabilities have already been demonstrated in multiple programmable quantum simulators consisting of $N \geq 50$ qubits based on trapped neutral atoms and ions, or superconducting qubits [12,180,28,74].

As an example, we present a feasible protocol for near-term implementation of our exact cluster model QCNN circuit via neutral Rydberg atoms [12,94], where long-range dipolar interactions allow high fidelity entangling gates [104] among distant qubits in a variable geometric arrangement. The qubits can be encoded in the hyperfine ground states, where one of the states can be coupled to the Rydberg level to perform efficient entangling operations via the Rydberg-blockade mechanism [104]; an explicit implementation scheme for every gate in Figure 4.2b is provided in Appendix A.7. Our QCNN at depth d with N input qubits requires at most $\frac{7N}{2}(1 - 3^{1-d}) + N3^{1-d}$ multi-qubit operations and $4d$ single-qubit rotations. For a realistic effective coupling strength $\Omega \sim 2\pi \times 10 - 100$ MHz and single-qubit coherence time $\tau \sim 200\,\mu$s limited by the Rydberg state lifetime, approximately $\Omega\tau \sim 2\pi \times 10^3 - 10^4$ multi-qubit operations can be performed, and a $d = 4$ QCNN on $N \sim 100$ qubits is feasible. These estimates are reasonably conservative as we have not considered advanced control techniques such as pulse-shaping [60], or potentially parallelizing independent multi-qubit operations.

Nothing is impossible, the word itself says "I'm possible"!

Audrey Hepburn

5

Enhancing Detection of Topological Order by Local Error Correction

The exploration of topologically-ordered states of matter is a long-standing goal at the interface of several subfields of the physical sciences. Such states feature intriguing physical properties such as long-range entanglement, emergent gauge fields and non-local correlations, and can aid

in realization of scalable fault-tolerant quantum computation. However, these same features also make creation, detection, and characterization of topologically-ordered states particularly challenging. Motivated by recent experimental demonstrations, we introduce a new approach for quantifying topological states—locally error-corrected decoration (LED)—by combining methods of error correction with ideas of renormalization-group flow. Our approach allows for efficient and robust identification of topological order, and is applicable in the presence of incoherent noise sources, making it particularly suitable for realistic experiments. We demonstrate the power of LED using numerical simulations of the toric code under a variety of perturbations, and we subsequently apply it to an experimental realization of a quantum spin liquid using a Rydberg-atom quantum simulator. Extensions to the characterization of other exotic states of matter are discussed.

5.1 Introduction and Motivations

Topological order is an exotic state of matter, which can occur when quantum fluctuations and local constraints stabilize a state with long-range entanglement [169]. With their non-local correlations, topologically-ordered states feature remarkable properties and can be used for protecting quantum information non-locally [169,121,155]. Yet, because these states appear to be liquid-like at short length-scales [134], they cannot be identified or characterized using any local order parameters. Instead, the canonical approach to discern topological order is to measure operators supported on large closed loops, the Wilson loops [75,171,169,69]. However, such operators are often challenging to identify or measure: while they have simple forms in certain fixed-point models, this is generally not the case for states realized experimentally. This is because experimental systems are typically affected by coherent perturbations generated by realistic Hamiltonians, or incoherent noise due to coupling with the environment (e.g., single qubit-flips or spontaneous emission). In these cases, the expectation values of the simple or 'bare' Wilson loop operators described above decay exponentially with the length of

the loop, which hinders the experimental certification of topological order.

To address these challenges, several methods have been developed to construct 'fattened' Wilson loops which do not decay with loop size. These include a systematic method utilizing quasi-adiabatic connections to the fixed-point models[75], as well as variational and tensor-network-based approaches[24,82,51,85]. Nevertheless, these methods are challenging to apply in realistic experiments, especially in the presence of decoherence and noise. Other signatures, such as topological entanglement entropy[91,102] are likewise difficult to measure in large systems.

Motivated by these considerations, we introduce a systematic method to construct and efficiently measure 'decorated' Wilson loop operators, a variant of the fattened loop operators. Our method, *locally error-corrected decoration (LED)*, leverages the error-correcting properties of the topological phase[155,92,48], and is also applicable to mixed states with incoherent noise. With a distinctive, hierarchical structure (Figure 5.1b), LED mimics the classification of quantum phases using RG flow[139,41]. For $\mathbb{Z}_2$ topological order, LED can be performed efficiently using classical post-processing of experimental measurements in a few fixed bases. These properties make LED applicable to a wide range of experiments where the prepared state is known to approximate a fixed-point state with zero correlation length (see Appendix B.7), and it allows for the verification of topological order at large length-scales, which is particularly challenging or impossible using conventional methods. In what follows, we demonstrate the power of the method by simulating perturbed toric code states generated using an efficient 2D tensor network sampling algorithm, and subsequently apply it to analyze the topologically-ordered states created in recent Rydberg atom array experiments[143].

5.2 LED Approach

The key idea of LED can be understood by considering Kitaev's toric code model, a canonical example of topological order. Specifically, we consider qubits localized on the edges of a square lattice,

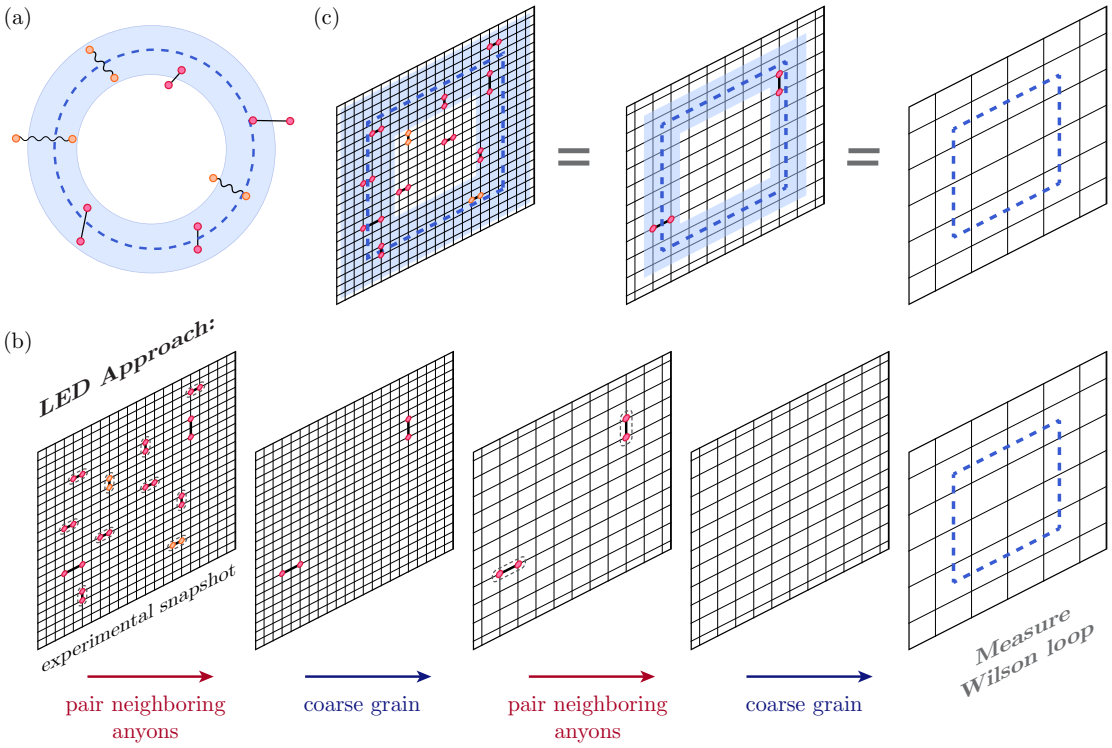


Figure 5.1: Detecting topological phases via LED. (a) In the absence of perturbations, a topologically-ordered state with zero correlation length such as Kitaev's toric code state[92] is characterized by $+1$ expectation values of 'bare' Wilson loop operators, which are typically tensor products of single-site operators (dotted blue loop). In realistic systems, however, coherent perturbations give rise to virtual anyon pairs (red dots/straight lines), and incoherent errors introduce physical anyon pairs (orange dots/wavy lines); this causes the expectation value of bare Wilson operators to decay exponentially with the loop's perimeter. To account for these local fluctuations, one can measure 'fattened' Wilson operators supported on an annulus (blue); the LED loops constitute one realization of this. (b) LED method to measure decorated Wilson loop observables for $\mathbb{Z}_2$ topological order: given an experimental snapshot of all qubits in the Z or X basis, one can obtain values for all stabilizer operators in that basis, thereby identifying the locations of all e or m anyons, respectively. Here, the qubits live on the links of the square lattice, and stabilizers are associated with vertices. In the first step, neighboring anyons are paired using a local decoder (dashed pairings), and each pair is removed by flipping the value(s) of qubit(s) lying on a path of minimal length connecting the two anyons; subsequently, the lattice is coarse-grained so that only a fraction of the original qubits remain. These two steps are iterated n times (here, $n = 2$), after which a bare Wilson loop is evaluated on the final, coarse-grained state. (c) The final, bare Wilson loop operator evaluated on the final state is equivalent to decorated Wilson loop operators evaluated at earlier iterations.

and the ideal, fixed-point Hamiltonian is defined as (see Section 2.2 and Ref. [92]):

$$H_{\mathrm{TC}} = -J\sum_{v} A_v - J\sum_{p} B_p \tag{5.1}$$

where $A_v = \prod_{i\in \mathrm{adj}(v)} X_i$, $B_p = \prod_{i\in \mathrm{adj}(p)} Z_i$, and $\mathrm{adj}(v)$ (resp., $\mathrm{adj}(p)$) denote the set of edges touching a given vertex v (plaquette p) of the lattice. The ground state space, given by the simultaneous $+1$ eigenspace of all *stabilizer operators* $\{A_v, B_p\}$, forms a quantum error-correcting code: all local operators either act trivially on ground states or couple them to excited states [92]. By measuring the stabilizers, one can detect the presence of excitations and apply a recovery procedure to return the system back to this ground state space.

In this model, contractible Wilson loops can be constructed by multiplying stabilizers A_v (B_p), so their expectation values in any ground state of H_{TC} are $+1$, independent of loop size. However, in realistic situations, the prepared state differs from the fixed-point state by local fluctuations such as coherent perturbations and incoherent errors (Figure 5.1a). This causes bare Wilson loops to decay exponentially with the number of locations where a fluctuation can intersect the loop (i.e., its perimeter).

Our LED approach begins with a measurement of all qubits in the same (Pauli-Z or Pauli-X) basis. For each measurement snapshot, one can calculate the stabilizer and Wilson loop values. Local fluctuations appear as stabilizer violations, which are identified with anyonic excitations [92] (Figure 5.1b). A local decoder partially removes such fluctuations by flipping measured qubits using only nearby stabilizer values. The simplest such local decoder can remove single-qubit errors, by flipping a qubit if and only if both adjacent vertices (resp., plaquettes) are occupied by an m (e-anyon). However, it cannot remove higher-weight errors, which flip two or more adjacent qubits. After correction, the lattice is coarse-grained, which can be also done efficiently on measurement snapshots (see Appendix B.2). Together, the anyon-pairing and coarse-graining steps are repeated for n layers.

Crucially, the weight of uncorrected errors is reduced in each layer, so that all local errors eventually become single-qubit errors which can be corrected by the decoder, as shown in Section 5.5; this mimics a real-space RG flow towards the fluctuation-free fixed-point state (see Appendix B.7). Finally, a bare Wilson loop is measured on the final, corrected and coarse-grained state (Figure 5.1b).

This bare Wilson loop operator measured on the final state is equivalent to a decorated Wilson loop operator measured on the original state (see Appendix B.2). In particular, this operator is determined solely by the fixed-point state and is independent of the specific fluctuations in the relevant system; this crucially differentiates LED from prior approaches to construct fattened loop operators[75,24,82] (see also Section 5.5). Moreover, all steps in LED can be performed in post-processing (see Appendix B.2), making LED uniquely suited for integration into experimental measurement procedures.

The hierarchical LED procedure is also inspired by the quantum convolutional neural network (QCNN) approach to phase classification, and the decorated Wilson loop operators resembles the "multiscale string order parameter" studied in Chapter 4. However, in this context, the LED framework is more general, as the hierarchical procedure shown in Figure 5.1 is merely one way of constructing LED Wilson operators of a particular diameter L and correction distance d (Figure 5.2a) [*]. More generally, given a desired value of d, one can construct the associated LED Wilson operators by choosing a local decoder which pairs anyons up to distance d (see Appendix B.3). The construction of Figure 5.1b, with alternating local-decoding and coarse-graining layers, is a particularly efficient way to construct local decoders and LED loops with longer-range (e.g. $d, L \propto 2^n$).

[*]Note that technically, the correction distance d is related to the annulus thickness cd by a decoder-dependent constant $c > 1$, since the range of information propagation is generally larger than the range of allowed anyon-pairings (see Appendix B.3).

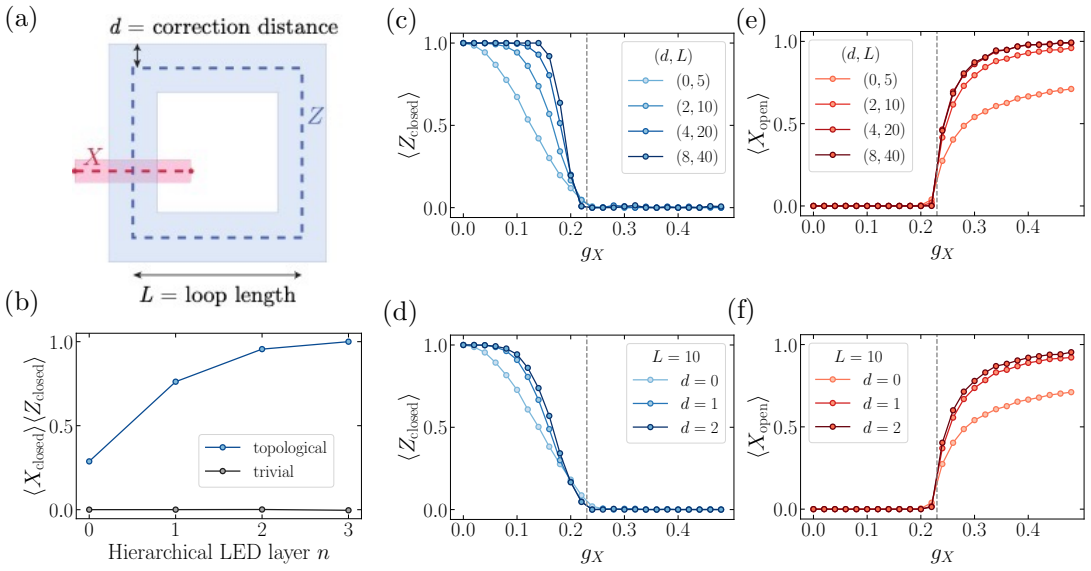


Figure 5.2: Numerical demonstration with coherently perturbed toric code states. (a) In a general construction of our LED Wilson loop operator, we use a local decoder which pairs anyons within a region of radius d (blue annulus). Conjugate LED open string operators (red stripe) anti-commute with Wilson loops, and hence must vanish in the topological phase. (b) Order parameter $\langle X_{\text{loop}}\rangle\langle Z_{\text{loop}}\rangle$ for a trivial state ($g_Z = 0.0, g_X = 0.26$) and a topological state ($g_Z = 0.12, g_X = 0.12$), upon varying n, using a distance-four patch decoder and coarse-graining blocksize two respectively (see Appendix B.3). (c) Output at different n along the $g_Z = 0.14$ line of the phase diagram. Gray dotted line is conjectured phase transition region. (d) Expectation values of generic LED Wilson loops with the same diameter L, using the pairing decoder ($d = 1$) and distance-four patch decoder ($d = 2$) without coarse-graining. (e,f) Corresponding expectation values of bare and decorated open string operators.

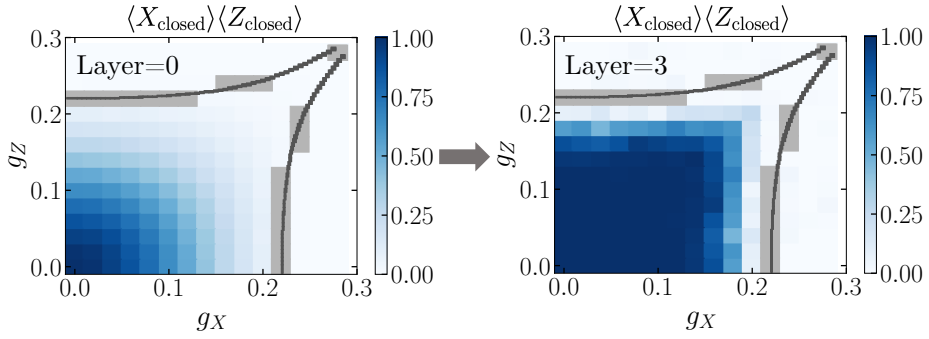


Figure 5.3: Two-dimensional phase diagrams for coherently perturbed toric code states. Order parameter values constructed from bare Wilson loops ($n = 0$) and LED Wilson loops ($n = 3$), using the same LED procedure as (c,e), across varying values of (g_X, g_Z). Dark gray regions are numerical estimates for the phase boundary between topological and trivial (see Appendix B.9). Light gray regions correspond to locations where sampling is expensive due to large correlation length.

5.3 Numerical Detection of Topological Order with Coherent Perturbations

To demonstrate the applicability of LED for coherent local perturbations to H_{TC}, we consider a family of states

$$|\psi(g_X, g_Z)\rangle = \frac{1}{\mathcal{N}} e^{g_X \sum_i X_i + g_Z \sum_i Z_i} |\psi_{\mathrm{TC}}\rangle \,, \tag{5.2}$$

generated by imaginary-time evolution of a toric code ground state $|\psi_{\mathrm{TC}}\rangle$. As each operator Z_i (resp., X_i) creates a pair of m anyons (e anyons), $|\psi(g_X, g_Z)\rangle$ contains virtual anyon fluctuations on top of $|\psi_{\mathrm{TC}}\rangle$. In the special case where $g_X = 0$, topological order is known to survive for perturbations $g_Z \leq g_c = 0.220343$, beyond which the m-anyons condense, driving a second-order phase transition into the Z paramagnet state [34]. More generally, $|\psi(g_X \neq 0, g_Z \neq 0)\rangle$ is also topologically-ordered for small g_X and g_Z, but the transitions to paramagnetic phases can occur at points which differ from g_c.

In testing LED, we numerically simulate projective measurements of $|\psi(g_X, g_Z)\rangle$ and use them as the input "experimental snapshot" in Figure 5.1b (see Appendix B.1). Figure 5.2b shows the value of the LED order parameter for a trivial and a topological state with $g_Z = 0.14$, when n is varied (and $d, L \propto 2^n$). Clearly, the order parameter stays at 0 for the trivial state, but increases from a small, finite value to one for the topological state as n is increased. Similar behavior is also observed throughout a one-dimensional parameter space in Figures 5.2c and 5.2d, whenever the correction distance d is increased, while keeping $d \ll L$ to prevent overcorrection (see Figure 5.7). Importantly, amplification occurs only if the input state is topological, and the order parameter approaches 0 for all trivial states.

Another important set of observables for characterizing topological order are X and Z open

string operators, which detect the transition from the topological phase to the trivial, paramagnet phase. Because LED Wilson Z-loop operators (resp., X-loop operators) are linear combinations of Z (X) closed loops supported on an annulus, they anti-commute with conjugate X (Z) open strings connecting the interior and exterior of the annulus. As such, the expectation value of long, open strings must flow to zero in the topological phase, where closed-loop LED operators flow to unity with increasing d [†]. The topological-to-trivial phase transition occurs when long, open X or Z strings acquire non-zero expectation value, due to the condensation of m or e anyons, respectively. Indeed, deep in the paramagnetic phase the state $\lim_{g_x \to \infty} |(g_x, g_z)\rangle$ is polarized along the X direction, and X open strings become unity. However, at generic points in the trivial phase, open strings also decay exponentially with length, due to local fluctuations of the opposite type (e or m anyons, respectively); nevertheless, LED can still amplify the trivial order signaled by open strings by removing the effect of local fluctuations. This behavior is demonstrated in our numerical simulations: in Figure 5.2e,f, open string expectation values stay at 0 in the topological phase, but are amplified and saturate to a non-zero value in the trivial (paramagnetic) phase.

Let us note that the boundary dividing the states whose LED operators approach zero and one does not necessarily correspond to the topological phase boundary: in general, it depends on the choice of decoder and coarse-graining length-scale. For instance, this is observed in Figure 5.3, where closed loops are nearly one after $n = 3$ layers for a large region within, but not fully encompassing, the topological phase. Nevertheless, we show that any state with LED Wilson loops flowing to unity at large d and L is topologically-ordered—that is, LED gives rise to a new sufficient condition or *witness* for topological order, even if it is not always a necessary condition (see Section 5.6). Thus, because LED amplifies the contrast between trivial states and a large class of topological states, it lowers the sample complexity of detecting topological order. This means that a statistically significant, nonzero expectation value of the order parameter can be obtained using substantially fewer

[†]This holds for any LED open string.

experimental repetitions (see Chapter 4, Ref.[70], and Appendix B.5).

5.4 Effect of Incoherent Errors

We next demonstrate the application of LED in the presence of incoherent local noise such as spontaneous emission or dephasing, which commonly occur in experiments. Because local decoders can recover topologically encoded information in the presence of small, local error channels[48,91] it is reasonable to ask whether the mixed states prepared in these systems exhibit topological ordering.

To study such examples, we introduce incoherent bit- and phase-flip errors by independently flipping, with probability p_{flip}, each measured qubit in a snapshot of $|\psi(g_X, g_Z)\rangle$. Here, we associate topological order with states that can be transformed into a ground state of H_{TC} via local operations. Our analysis then suggests that the resulting mixed-state phase space contains a $\mathbb{Z}_2$-topological phase, a Z-paramagnet, an X-paramagnet, and a disordered phase with large incoherent error rates. However, it is especially difficult to distinguish the topological and disordered phases using measurements of bare operators alone: in both phases, open strings remain close to zero, while bare Wilson loops decay exponentially with perimeter as $e^{-\alpha L}$, where the exponent α interpolates smoothly between the phases (Figure 5.4a,b). This is in contrast to the paramagnet phases, where closed loops exhibit similar behavior, but certain open strings decay with the same exponent α as the closed loops[59].

However, upon studying the behavior of LED operators, one finds that the mixed-state phase space exhibits two qualitatively different regimes (Figure 5.4b). Upon increasing d, LED reduces α to 0 in the 'correctable' regime, while α grows in the 'uncorrectable' regime. Further, correctable states with small p_{flip} are connected to topologically-ordered pure states, suggesting these mixed states are topologically-ordered as well. Indeed, we show theoretically that any state in the correctable regime has long-range topological order at least up to distances of order $L - d$ (see Sec-

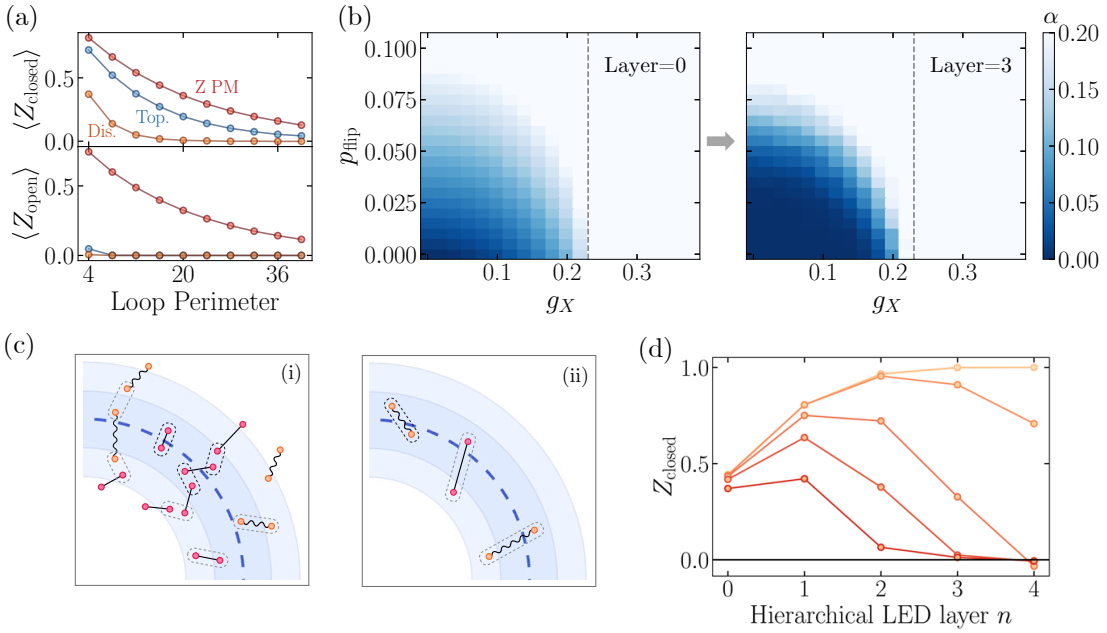


Figure 5.4: Application to mixed states. (a) Without error correction, generic points in the topological and trivial, disordered phase ($g_Z = 0.12, g_X = 0.18, p_{\text{flip}} = 0.0$ and $g_Z = 0.06, g_X = 0.0, p_{\text{flip}} = 0.11$ resp. shown in the plot) appear very similar qualitatively, as closed loops decay exponentially with loop perimeter in both cases, while open strings remain close to zero. In contrast, in the trivial, paramagnet phase ($g_Z = 0.32, g_X = 0.2, p_{\text{flip}} = 0$), open loops decay with the same perimeter-law as closed loops. (b) $g_Z = 0.14$ slice of mixed state phase diagram, containing topological, disordered, and X-paramagnetic phases. These phases are associated with fixed-point states $g_x = g_z = p_{\text{flip}} = 0, g_z \to \infty, g_x \to \infty$, and $p_{\text{flip}} \to 0.5$, respectively. The flow of the closed-loop decay exponent α under LED provides a sharp divider between two kinds of perimeter-law decay, observed in different regimes of the mixed-state phase diagram. (c) In the uncorrectable regime (i), the local decoder of LED pairs anyons incorrectly, resulting in perimeter-law decay with large α in disordered and paramagnet phases. Moreover, the probability of such an incorrect pairing can increase with the number n of LED iterations. Here, the black pairings are made by LED at or before one specific value of n, and gray pairings are made upon performing one additional LED iteration. In the correctable (topological) regime (ii), increasing n can reduce α to zero, as fluctuations of higher characteristic length ξ can be reliably corrected using only local information. In the conceptual framework where an LED operator is embedded in a surface code on an annulus (see Figure 5.2a), incorrect pairings corresponds to logical errors (e.g. X_L). (d) Expectation values of LED loop observables upon increasing n ($d, L \propto 2^n$), in thermal states of varying temperatures (between 0 and 0.35, with darker colors indicating higher temperatures) and $p_{\text{flip}} = 0.02$.

tion 5.6). Thus, LED enables a more efficient and accurate detection of topological order.

The ability of LED to distinguish between the topological and disordered phases can be understood by analogy to quantum error correction. Conceptually, since any given LED loop operator is supported on an annulus, we can consider this operator as being embedded in a surface code on this annulus with open boundary conditions, which supports a logical qubit. For example, in considering an LED Z-loop as in Figure 5.2a, the LED loop operator corresponds to the logical-Z operator for this qubit after applying error correction. Meanwhile, any X-string connecting the interior of the annulus to the exterior corresponds to a logical-X operator, as it anti-commutes with the logical-Z operator. In this framework, the decay rate α of Wilson loops corresponds to a local logical error rate per unit length, and in the correctable phase, LED-based decoding succeeds with high probability as long as the code distance d is sufficiently large (Figure 5.4c). However, in the uncorrectable phase, such as when p_{flip} is above the error correction threshold or when long, open strings condense in a paramagnetic phase, decoding cannot correctly pair anyons, resulting in a high rate of logical errors[48].

The above results are deeply rooted in the stability of topological order against local perturbations. In contrast, any finite temperature destroys long-range topological order as it leads to freely propagating thermal anyons. In Figure 5.4d, we consider the toric code model at finite temperature, with local incoherent errors, and find that the LED loop operators indeed approach zero upon increasing n. Interestingly, their expectation values flow non-monotonically: the loops are amplified at small n before eventually turning to 0. This occurs because of a competition between two effects: thermal anyons are uncorrectable, so their density accumulates under RG flow; however, local fluctuations are corrected at early layers, which initially amplifies LED loop expectation values. Because loops at different n probe correlations at different length-scales, the turning point in these curves can be used to identify the characteristic length-scale of separation between thermal anyons, or equivalently, the system's temperature.

5.5 Arbitrary Local Perturbations

We now prove a key property of LED: any state which differs from the fixed-point state by an arbitrary local perturbation has LED Wilson loops which flow to one. This property implies that LED loop operators are independent of the exact perturbation, unlike the fattened Wilson loops of Refs.[75,102]. For concreteness, we examine perturbations on top of a toric code ground state.

To prove our claim, we consider a local unitary operator $\mathcal{O}$ supported on a local region A of diameter l. It follows that $\mathcal{O}$ can only flip a stabilizer from $+1$ to -1 if it overlaps with A, and $\mathcal{O}$ cannot couple any ground state $|\psi\rangle$ to another ground state $|\psi'\rangle \neq |\psi\rangle$. We now show that LED removes all flipped stabilizers after $1+\log_b d$ layers, where b is the coarse-graining length-scale. Because the coarse-graining step effectively reduces $l \to l/b$, after $\log_b d$ layers, there are three possibilities:

1. A becomes fully contained within a single $b \times b$ region at some layer $c < \log_b d$. Then, A has zero support after another layer of coarse-graining and disappears.

2. Before iteration $\log_b d$, $\mathcal{O}$ is supported on two adjacent $b \times b$ regions. Then, $\mathcal{O}$ becomes a single-qubit error after this iteration and is removed by the subsequent LED step.

3. Before iteration $\log_b d$, $\mathcal{O}$ is supported at the corner of three or four regions. In this case, it becomes a two-qubit diagonal error after this iteration, which can also be removed by the subsequent LED step.

Notice that handling the third case requires the inclusion of diagonal pairing in the pairing decoder. Finally, while this proof focuses on the pairing decoder, it also generalizes directly to more advanced local decoders, such as the patch-based decoder presented in Appendix B.3, when they are combined with coarse-graining.

5.6 Topological Order Witness

In this section, we show that LED provides a topological order witness—that is, it does not misclassify any trivial product state as topological. For simplicity, we study the case of $\mathbb{Z}_2$ topological order on a surface with trivial topology, where the fixed-point state is the unique ground state $|\psi_{\mathrm{TC}}\rangle$ of H_{TC}. We begin by considering the ideal case where LED operators flow to unity.

Theorem 1. *Let $|\psi\rangle$ be an arbitrary input state defined on a surface with trivial topology. Then, after performing LED with correction distance d, assume the resultant state $|\psi_d\rangle$ has, as a subsystem, qubits living on the links of a square lattice, as in the toric code. Then, if the stabilizer expectation values $\left\langle \frac{1+A_v}{2} \right\rangle = \left\langle \frac{1+B_p}{2} \right\rangle = 1$ at every vertex v and plaquette p of the subsystem, then, the input state $|\psi\rangle$ is topologically-ordered, in the sense that it is connected to an output state of the form $|\psi_d\rangle = |\psi_{\mathrm{TC}}\rangle \otimes |\varphi_{\mathrm{anc}}\rangle$ by generalized local unitary (gLU) transformation of depth $O(d)$.*

Proof. The LED procedure forms a local quantum channel, and we begin our proof by constructing a purification of this channel. To mediate stabilizer measurement and local error correction, one can first introduce an ancilla in the state $|0\rangle$ at every vertex and plaquette. Next, a sequence of Hadamard and controlled-NOT (CNOT) gates is applied such that a Z-basis measurement on an ancilla is equivalent to the associated stabilizer measurement of A_v or B_p. Then, local quantum error correction is performed using a local unitary evolution on the combined system, which contains the original state and the added ancilla qubits. This local unitary evolution applies gates which perform X and Z spin flips on the system qubits, conditioned on the state of the ancilla qubits. Finally, the coarse-graining step can also be performed with local unitary transformations by using a quantum circuit corresponding to a multiscale entanglement renormalization ansatz (MERA) representation of the fixed-point state [1]. The transformations generated by introducing product state ancillas and performing local unitary operations are called generalized local unitaries (gLU); this class of transformations includes our LED procedure described above and is known to preserve phase boundaries [40].

If the system part of the final state, $|\psi_d\rangle \in \mathcal{H}_{\text{sys}} \otimes \mathcal{H}_{\text{anc}}$, has stabilizer expectation values $\left\langle \frac{1+A_v}{2} \right\rangle = \left\langle \frac{1+B_p}{2} \right\rangle = 1$ at every vertex v and plaquette p, then $|\psi_d\rangle$ must belong to the ground state space of the toric code. This is because the projector onto the ground state space is given by the product of all the stabilizers $P_{\text{TC}} = \prod_v \frac{1+A_v}{2} \prod_p \frac{1+B_p}{2}$. On a surface with trivial topology, there is a unique state $|\psi_{\text{TC}}\rangle$, so the output state factors into $|\psi_d\rangle = |\psi_{\text{TC}}\rangle \otimes |\varphi_{\text{anc}}\rangle$. □

If, in addition to the conditions in Theorem 1, we further assume the output ancillas $|\varphi_{\text{anc}}\rangle$ are in a trivial state, the above proof guarantees that the input state is in the toric code phase. However, we do not certify this condition holds, which is in general more difficult: measurements in multiple bases are needed to uniquely determine $|\varphi_{\text{anc}}\rangle$, and with incoherent errors, the ancillas are described by a mixed state. Instead, LED certifies that the toric code state can be *distilled* from the input state by gLU transformations. Because long-range entanglement cannot be created from a trivial state by gLU transformations[40], Theorem 1 implies that LED operators flowing to unity forms a sufficient condition for topological order, or equivalently, a topological order witness (see also Ref.[69]).

While the above argument works well in theory, any practical system cannot measure LED observables equal to one with infinite precision. Indeed, even infinitesimal local perturbations to the toric code ground state, such as $e^{-i\varepsilon H}|\psi_{\text{TC}}\rangle$ for arbitrarily small ε and some local Hamiltonian H, can create error strings larger than the correction length d. This causes LED loop expectation values to decay exponentially, even in the topological phase. To show that LED still provides a topological order witness in the presence of local perturbations, finite measurement errors, and finite system size, we show the following Theorem:

Theorem 2. *Consider an arbitrary state $|\psi\rangle$ and LED with correction distance d, as in Theorem 1. Suppose the corresponding subsystem of $|\psi_d\rangle$ has stabilizer expectation values $\left\langle \frac{1+A_v}{2} \right\rangle > 1-\varepsilon$, $\left\langle \frac{1+B_p}{2} \right\rangle > 1-\varepsilon$ at every vertex v and plaquette p. Then, the input state $|\psi\rangle$ exhibits topological ordering at least up to a length-scale $O(\mathscr{L}-d)$; that is, the input state $|\psi\rangle$ cannot be prepared using a*

local quantum circuit of depth less than $O(\mathscr{L} - d)$, where $\mathscr{L} \sim 1/\sqrt{\varepsilon}$.

Our proof of Theorem 2 hinges on the following two Lemmas:

Lemma 3. *Given an output state $|\psi_d\rangle$ satisfying the conditions of Theorem 2, and a simply connected $(\mathscr{L} - 2) \times (\mathscr{L} - 2)$ square region R on the system part, the reduced density matrix $\rho_d = \mathrm{Tr}_{R^c}[|\psi_d\rangle \langle\psi_d|]$ is indistinguishable from the toric code reduced density matrix $\sigma_{\mathrm{TC}} = \mathrm{Tr}_{R^c}[|\psi_{\mathrm{TC}}\rangle \langle\psi_{\mathrm{TC}}|]$ defined on the same region, up to the bound $||\rho_d - \sigma_{\mathrm{TC}}|| \leq \max\left(\sqrt{\varepsilon}, 2\mathscr{L}^2\varepsilon\right)$*

Proof. To bound the trace distance, we will use the fact that our state ρ_d locally looks almost the same as the toric code state. Specifically, trace distance is related to distinguishability by [122]

$$||\rho - \sigma|| = \frac{1}{2} \sup_{||O|| \leq 1} \mathrm{Tr}\left[O(\rho - \sigma)\right]. \tag{5.3}$$

To upper bound this, we can consider all possible unit norm operators O. Specifically, O can always be written as a linear combination of Pauli strings. These Pauli strings can be analyzed by considering two cases, closed strings and open strings. First consider operators C supported on R, which commute with all stabilizers A_v, B_p supported on a slightly larger $\mathscr{L} \times \mathscr{L}$ region (1-ball or R) , constructed by expanding on all sides by one unit cell. These operators must be products of contractible Wilson loops, and hence can be written as product of stabilizers. Therefore, $\mathrm{Tr}[\sigma_{\mathrm{TC}} C] = 1$ for the toric code state. For the LED output state, we instead bound the expectation value of $P_{\mathrm{TC}} = \prod_v \frac{1+A_v}{2} \prod_p \frac{1+B_p}{2}$, where the product over v (resp., p) runs over all vertices (plaquettes) within the $\mathscr{L} \times \mathscr{L}$ region. The expectation value of this projector, evaluated on the output state, is given by $\langle\psi_d|P_{\mathrm{TC}}|\psi_d\rangle$. To lower bound this quantity, we first notice that every term in P_{TC} has spectrum $\{0, 1\}$, and that the terms are mutually commuting.

This allows us to approximate $\langle P_{\mathrm{TC}}\rangle$ using the individual expectation values $\left\langle \frac{1+A_v}{2} \right\rangle > 1 - \varepsilon$ and $\left\langle \frac{1+B_p}{2} \right\rangle > 1 - \varepsilon$. To do so, we note that if two commuting operators A and B, each with spectrum

$\{0,1\}$, satisfy $\langle A\rangle > 1-\varepsilon$ and $\langle B\rangle > 1-\varepsilon$, then $\langle AB\rangle > 1-2\varepsilon$. To show this, we add the two individual bounds to obtain $\langle A\rangle + \langle B\rangle > 2-2\varepsilon$; moreover, since AB has spectrum $\{0,1\}$, we have $\langle A\rangle + \langle B\rangle - \langle AB\rangle < 1$. It thus follows that $\langle AB\rangle > 1-2\varepsilon$, and upon applying this recursively to include all vertex and plaquette terms within the $\mathscr{L}\times\mathscr{L}$ region, we obtain $\langle\psi_d|P_{\mathrm{TC}}|\psi_d\rangle > 1-2\mathscr{L}^2\varepsilon$.

Using this fact, along with $CP_{\mathrm{TC}} = P_{\mathrm{TC}}$ we can similarly lower bound the expectation value $\mathrm{Tr}[\rho_d C] \geq 1-4\mathcal{L}^2\varepsilon$. Thus, for any C, we have $\frac{1}{2}|\mathrm{Tr}\left[C(\rho_d-\sigma_{\mathrm{TC}})\right]| \leq 2\mathcal{L}^2\varepsilon$.

Next, we consider operators O which anti-commute with some stabilizers. In particular, these operators are Pauli strings with at least one endpoint (open strings). Naturally, their expectation value vanishes in the toric code state. To see this, let S be one of the stabilizers which anti-commutes with O. Now, S and O form an anti-commuting pair of Pauli operators, so they satisfy an uncertainty relation $\langle S\rangle^2 + \langle O\rangle^2 \leq 1$. As such, in the toric code where $\langle S\rangle = 1$, this implies $\langle O\rangle = 0$. Similarly, the condition $\mathrm{Tr}[\rho_d S] \geq 1-2\varepsilon$ leads to the upper bound $\mathrm{Tr}[\rho_d O] \leq 2\sqrt{\varepsilon}$. Combining these two results, we see the trace distance is at most $T(\rho_d,\sigma_{\mathrm{TC}}) \leq \max\left(\sqrt{\varepsilon}, 2\mathscr{L}^2\varepsilon\right)$. □

Lemma 4. *Consider an input state $|\psi\rangle$ and an LED procedure satisfying the conditions of Theorem 2. Then the final state $|\psi_d\rangle$ after LED cannot be prepared using a local quantum circuit with depth less than $O(\mathscr{L}) \sim O(1/\sqrt{\varepsilon})$.*

To prove these results, we extend and generalize the proof techniques developed by Ref.[69].

Proof. Consider a pair of Z-basis and X-basis Wilson loops A and B, supported on two overlapping annuli. The twist product of the two operators $A\infty B$ is defined such that at one intersection region, operator B is applied first, while at the other operator A is applied first (Fig. 5.5). By the arguments of Ref.[69], such a pair of locally non-commuting observables, whose twist product does not factorize into a product of the individual observables, can serve as a witness for long-range entanglement. In our case, since Z- and X-strings locally anti-commute, we can remove the twist to get $A\infty B = -AB$.

Then, Ref.[69] showed that, assuming the observables A and B satisfy an additional important property called local invisibility (see below), then the following correlation serves as a witness for long-range order

$$|\langle\psi| A\infty B |\psi\rangle - \langle\psi| A |\psi\rangle \langle\psi| B |\psi\rangle| > 0. \tag{5.4}$$

Specifically, the state $|\psi\rangle$ cannot be prepared from a trivial state by a circuit of depth $O(L)$, where L is the separation between the two intersection regions. Indeed, in the exact case, where the expectation value of large Wilson loops with $|\psi_d\rangle$ are one, these results can be directly applied.

However, the proofs in Ref.[69] do not immediately apply to the realistic case considered here, where stabilizers have expectation value $1-\varepsilon$, and residual entanglement between the ancilla and system qubits prevents exact knowledge of the state. Nevertheless, with sufficient care and a few additional assumptions, approximate versions of key results in Ref.[69] can be recovered.

First, we develop a notion of *approximate local invisibility*. Throughout, we follow the spirit of the proofs in Section III of Ref.[69]; the reader is encouraged to consult the original reference for additional details and insights.

Definition 5 (Approximate (Δ, r, t)-local invisibility.)**.** Let A be a region of radius r and B be a t-ball around A. An operator O with unit norm is (Δ, r, t)-locally invisible with respect to a state $|\psi\rangle$ if, for any state $|\varphi\rangle$ whose reduced density matrix on B is equivalent to $|\psi\rangle$, it satisfies

$$\left|\left| \frac{\mathrm{Tr}_{A^c}[O|\varphi\rangle\langle\varphi|O^\dagger]}{\mathrm{Tr}[O|\varphi\rangle\langle\varphi|O^\dagger]} - \mathrm{Tr}_{A^c}[|\psi\rangle\langle\psi|] \right|\right| \leq \Delta, \tag{5.5}$$

where the norm is the standard trace norm. In other words, locally invisible operators leave local reduced density matrices approximately unchanged. Note that we restrict to states $|\varphi\rangle$ for which the expectation value does not vanish, such that this remains well-defined. This subtlety is also present

in the original definition of Ref.[69].

Next, we will show that Wilson loops that nearly stabilize $|\psi_d\rangle$ are approximately locally invisible. Let A be a region of radius r, which can only cover a patch of the loop. Furthermore, let $t = 0$, i.e. region B is identical to region A. Since the Wilson loop is a tensor product of local unitaries, we can write $O = O_B \otimes O_{B^c}$ and $O_{B^c}^\dagger O_{B^c} = 1$. This allows us to work directly with the reduced density matrices of region B, and we can use Lemma 3 to reduce to the toric code case

$$\mathrm{Tr}_{A^c}[O\rho_d O] = \mathrm{Tr}_{A^c}[O\sigma_{\mathrm{TC}}O] + \mathrm{Tr}_{A^c}[O(\rho_d - \sigma_{\mathrm{TC}})O] \tag{5.6}$$

Indeed, since O is locally invisible with respect to to the toric code, this gives us our result, where the error term depends on the size of B.

$$\left|\left|\mathrm{Tr}_{A^c}\left[O\rho_d O\right] - \mathrm{Tr}_{A^c}\left[\sigma_{\mathrm{TC}}\right]\right|\right| \leq \max\left(\sqrt{\varepsilon}, 2(r+1)^2\varepsilon\right) \tag{5.7}$$

More microscopically, O spans the region, so locally looks like a logical operator. The reduced density matrix σ_{TC} on region B is an equal weight mixture of all logical states, so O leaves it invariant.

This shows that Wilson loops are (Δ, r, t)-locally invisible with respect to $|\psi_d\rangle$ for $t \geq 0$ and $\Delta = \max\left(\sqrt{\varepsilon}, 2(r+1)^2\varepsilon\right)$. When combined with the fact that Wilson loops have large expectation value on $|\psi_d\rangle$, this will serve as a witness for long-range topological order.

To prove this, we need to confirm that, even for the weaker notion of approximate local invisibility, the twist product approximately factorizes for trivial states.

Lemma 6. *The twist product of two (Δ, r, t) locally invisible operators A and B, acting on a trivial product state $|\psi\rangle = |00\ldots0\rangle$, must satisfy*

$$|\langle\psi| A \infty B |\psi\rangle - \langle\psi| A |\psi\rangle \langle\psi| B |\psi\rangle| \leq O(\sqrt{\Delta R/r}) \tag{5.8}$$

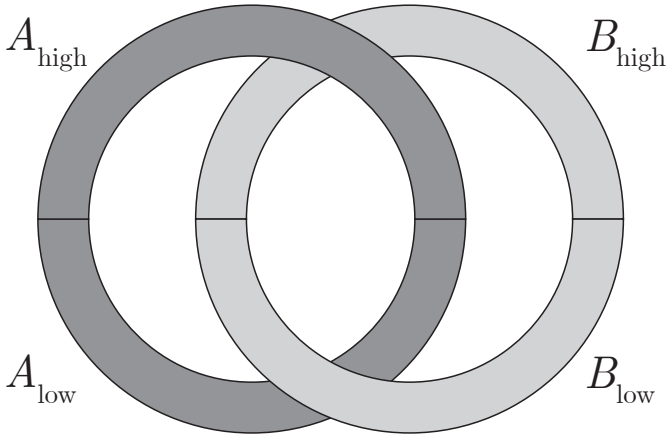


Figure 5.5: Consider two operators A and B, each supported on an annulus which intersects at two regions, one higher and one lower. The twist product $A \infty B$ consists of applying B first in the higher region, and A first in the lower region (see also Figure 1 and Equation (4) in Ref.[69]). In the diagram, the order of operations goes back to front. Furthermore, for tensor product operators which can be written as $A = A_{\text{low}} A_{\text{high}}$, $B = B_{\text{low}} B_{\text{high}}$, we can write the twist product as $A \infty B = B_{\text{low}} A_{\text{high}} A_{\text{low}} B_{\text{high}}$.

where A and B are supported on two annuli which intersect at two regions (Figure 5.5) whose separation is $\geq 2(r+t)$.

Proof. In the first step of the proof, we bound the expectation value of $\Pi_{\mathcal{R}} = \prod_{i \in \mathcal{R}} |0\rangle \langle 0|_i$, the projector onto $|\psi\rangle$ supported on region $\mathcal{R}$ evaluated with respect to $O|\psi\rangle$ for unitary O.

Invoking the definition of local invisibility, we show for A of radius r, that $\langle\psi| O^\dagger \Pi_A O |\psi\rangle \geq 1 - \Delta$.

$$\left|\left| \text{Tr}_{A^c} \left[O|\psi\rangle \langle\psi| O^\dagger \right] - \Pi_A \right|\right| \leq \Delta \tag{5.9}$$

Thus, the expectation value of the observable Π_A satisfies:

$$|\text{Tr}\left[\Pi_A O|\psi\rangle \langle\psi| O^\dagger\right] - 1| \leq \Delta \tag{5.10}$$

$$\langle\psi| O^\dagger \Pi_A O |\psi\rangle \geq 1 - \Delta \tag{5.11}$$

Next, we can use the fact that $\Pi_{\mathcal{R}}$ can be written as a product of Π_A approximately R/r times. It follows from the same union bound argument in Lemma 3, that $\langle\psi| O^\dagger \Pi_{\mathcal{R}} O |\psi\rangle \geq 1 - \Delta R/r$. This result can be used to bound the distance between the two states,

$$||\Pi_{\mathcal{R}} O |\psi\rangle - O |\psi\rangle ||^2 \leq \Delta R/r, \tag{5.12}$$

by noticing the left side is equal to $1 - \langle\psi| O\Pi_{\mathcal{R}} O |\psi\rangle$. We will use this below, to show that we can replace $O |\psi\rangle$ with $\Pi_{\mathcal{R}} O |\psi\rangle$ without incurring significant error.

To prove Lemma 4, we use the same construction as Ref.[69]. Specifically, we want to use the above result to show that $A \infty B |\psi\rangle = (A\Pi_{\mathcal{R}}) \infty B |\psi\rangle + O(\Delta \mathcal{L}^2)$. We do this by carefully inserting projectors. For Wilson loops, which are tensor product operators, we the twist product can be split up as follows (see Figure 5.5),

$$A \infty B = B_{\text{low}} A_{\text{high}} A_{\text{low}} B_{\text{high}}. \tag{5.13}$$

Our above result implies $\Pi_{\mathcal{R}} B_{\text{high}} |\psi\rangle = B_{\text{high}} |\psi\rangle + O(\Delta R/r)$. We can subsequently pull projectors from $|\psi\rangle$ to cover the region of support of A, e.g. low and the parts of high. Thus, we get

$$\begin{aligned} A \infty B |\psi\rangle =& B_{\text{low}} A_{\text{high}} \Pi_{\text{high}} A_{\text{low}} \Pi_{\text{low}} B_{\text{high}} |\psi\rangle \\ &+ O(\sqrt{\Delta R/r}) \\ =& (A\Pi) \infty B |\psi\rangle + O(\sqrt{\Delta R/r}) \end{aligned} \tag{5.14}$$

as we wanted. Finally, we this implies the expectation value of the twist product approximately factorizes $\langle\psi| A \infty B |\psi\rangle = \langle\psi| A |\psi\rangle \langle\psi| B |\psi\rangle + O(\sqrt{\Delta R/r})$. □

Although we proved that the twist product must factorize for the trivial product state, this holds for a much wider class of short-range entangled states, generated from a trivial state by a finite-depth

unitary circuit. In particular, using the same argument as Lemma III.3 of Ref.[69] shows that, given a (Δ, r, t)-locally invisible operator O and state $|\psi\rangle$, if we evolve under a local unitary circuit W of depth d, then the operator $WOW^\dagger$ is $(\Delta, r-d, t+2d)$-locally indistinguishable with respect to $W|\psi\rangle$. This will be used to show that non-factorizability of the twist product lower bounds the depth of a quantum circuit required to produce the state from the trivial product state.

Finally, we can use the trace distance Lemma 3, to show the twist product does not factorize for $|\psi_d\rangle$. Specifically, let A and B be Wilson loops supported on an $\mathscr{L} \times \mathscr{L}$ region. Then,

$$|\langle\psi_d| A \infty B |\psi_d\rangle - \langle\psi_d| A |\psi_d\rangle \langle\psi_d| B |\psi_d\rangle| \geq 2 - c\mathscr{L}^2\varepsilon \tag{5.15}$$

for a constant c. Combining this Lemma 6 for trivial states, we see the bound is violated when

$$2 - c\mathscr{L}^2\varepsilon > c'\sqrt{\Delta\mathscr{L}/r} \tag{5.16}$$

$$2 - c\mathscr{L}^2\varepsilon - c' 2(r+1)\sqrt{\mathscr{L}\varepsilon/r} \geq 0 \tag{5.17}$$

We choose r to be a constant fraction of $\mathscr{L}$, and see that we can roughly make $\mathscr{L} \sim O(1/\sqrt{\varepsilon})$ and still certify long-range order. In particular, the state $|\psi_d\rangle$ cannot be generated by a finite depth circuit of depth smaller than $r \sim O(\mathscr{L})$. Since $|\psi_d\rangle$ is connected to the input state by a depth-d quantum circuit, this implies LED Wilson loops close to one certify topological order up to length-scales $O(\mathscr{L} - d)$.

□

Upon combining the result of Lemma 4 with the fact that our LED procedure corresponds to a local quantum circuit with depth $O(d)$, we find that the original input state $|\psi\rangle$ cannot be prepared using a quantum circuit of depth smaller than $O(\mathscr{L} - d)$—which is precisely the statement of Theorem 2. So, if we measure loops of length $L \gg d$ to be $1-\varepsilon$, this shows that LED provides a

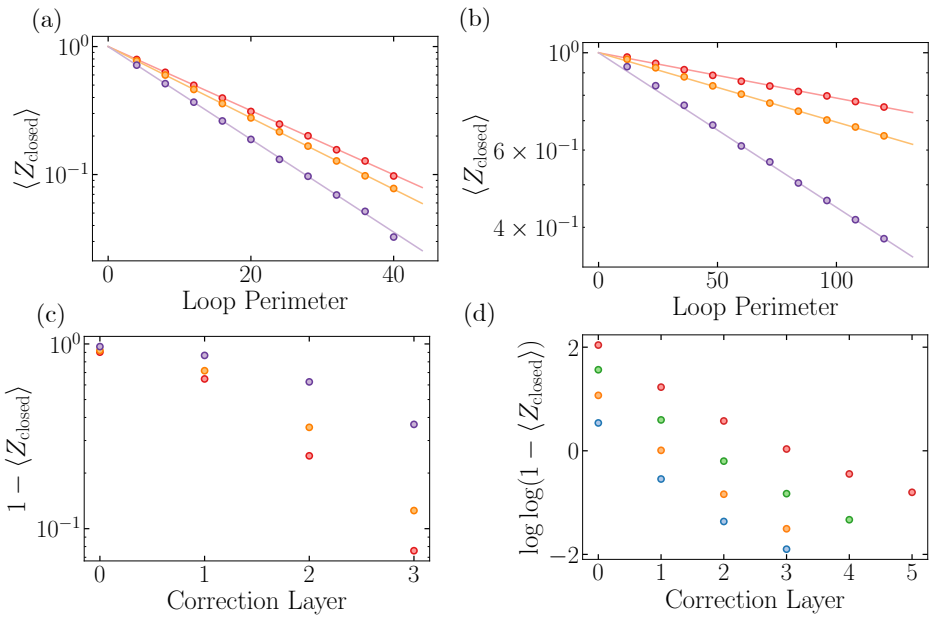


Figure 5.6: Perimeter-law decay of Wilson loops. The exponential decay with loop size is clearly visible at various points in the topological phase – $g_Z = 0, g_X = 0.18, p_{\text{flip}} = 0$ (orange), $g_Z = 0.18, g_X = 0.18, p_{\text{flip}} = 0$ (red), and $g_Z = 0.10, g_X = 0.18, p_{\text{flip}} = 0.03$ (purple). This is observed for both (a) uncorrected loops and (b) $d = 6$ corrected loops under two layers of $d = 3$ MWPM patch decoding. (c) LED Wilson loops appear to approach one faster than exponential in n. (d) In a model with only incoherent errors ($p_{\text{flip}} = 0.02$ (blue), 0.03 (orange), 0.04 (green), 0.05 (red)), we can study the effect of even more layers, where we see hints that the decay is doubly-exponential in n, or exponential in $d \sim 2^n$.

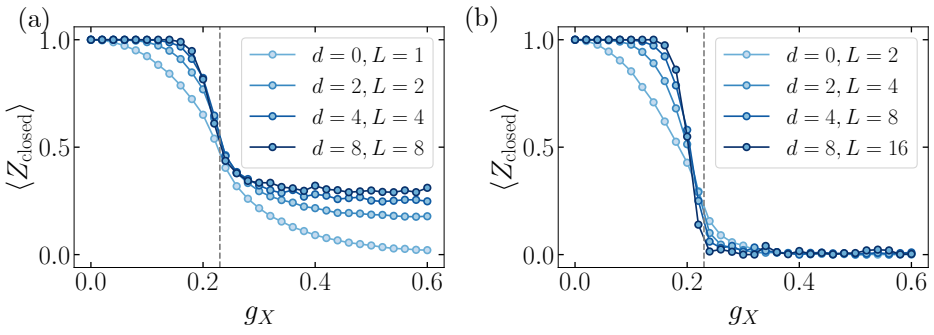


Figure 5.7: Over-correction for small Wilson loops. a) Small LED loops with $d > L$ can give non-zero signal deep in the trivial phase, as in this regime, correction can pump anyons from the interior of the annulus to the exterior. b) However, once $d < L$, the effects of over-correction become insignificant. LED Wilson loops are theoretically expected to certify topological order in the regime of $d \ll L$.

topological order witness up to length-scales of $O(L/\sqrt{\varepsilon})$.

We now discuss how these theoretical results are reflected in our numerical simulations. First, when fluctuations are local, the probability of having an error string of length ℓ decays exponentially with ℓ, and the exponent is determined by the characteristic length-scale ξ of fluctuations. In these systems, we expect the error rate after an optimal LED procedure with correction distance d to be given by $\varepsilon(d) \propto e^{-d/\xi}$, so the length-scale $\mathscr{L}$ with which we can certify topological order grows exponentially with d. Second, when LED uses the hierarchical, anyon-pairing decoder, the anyon density is observed to decrease faster than exponentially in the number n of LED steps (Figure 5.6). In this case, both the measured stabilizer size and the correction distance d grow exponentially with n, which implies that the certification length-scale $\mathscr{L}$ grows at least exponentially with n as well. Third, our argument does not certify topological order to any length-scale when $L < d$; this is because the support of such an LED operator no longer has an interior, potentially giving rise to signal even in the trivial phase. Indeed, this is reflected in our numerics as well (Figure 5.7).

Finally, we notice a connection between the flow of LED closed loops to $+1$ expectation value and the input state's topological entanglement entropy. Specifically, whenever the underlying state has finite correlation length and area-law entanglement, the LED loops which flow to unity are direct witnesses for nonzero topological entanglement entropy. This is because the flow to unity implies that the superselection sector of the region—that is, which anyons are contained in the region—is a well-defined quantum number. Indeed, the $+1$ expectation value of all contractible LED Wilson loops implies the region enclosed by any such loop is in the vacuum superselection sector. These constraints lead to a topological entanglement entropy of $2 \log 2$ [102,91]. As described in Ref. [172], this argument also generalizes when the decoder is non-local (see Appendix B.10).

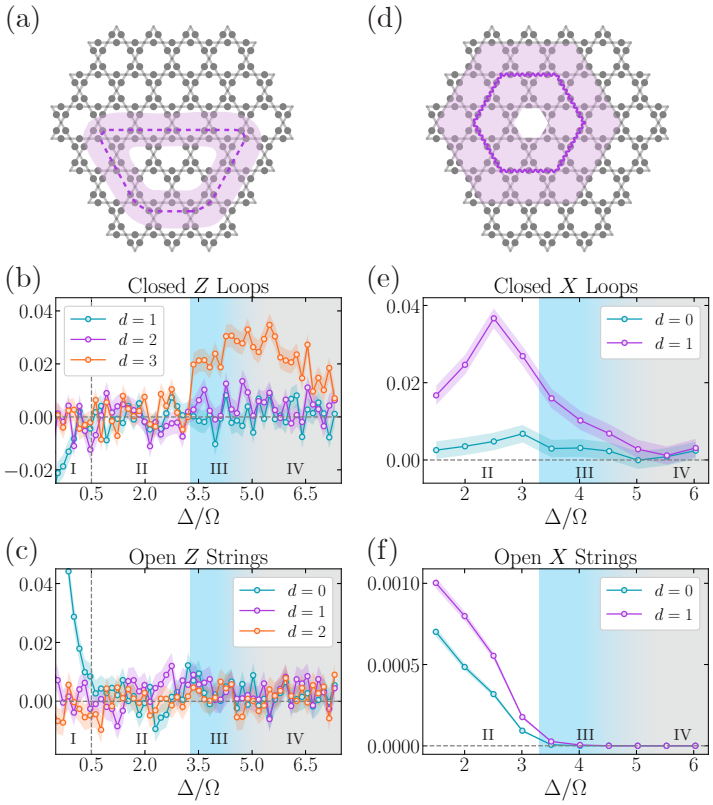


Figure 5.8: Enhancing experimental detection of $\mathbb{Z}_2$ spin liquid. (a,d) In the experiment[143], 219 qubits are placed on the links of a kagome lattice. Upon applying LED, the Z and X closed-loop observables are amplified for certain ranges of Δ/Ω. The shaded purple regions show the support of large, decorated Wilson loops with $d = 1$. (b,e) Expectation value of Wilson loops depicted in (a,c) for different correction layers n. Plotted error bars (shaded regions) show expected variation (one standard deviation) of the mean. The regime in which both types of loops are amplified corresponds to the spin-liquid regime identified in Ref.[143] (shaded blue region). (c,f) The behavior of expectation values of open Z- and X-strings under LED further confirms our findings, as both types of open strings stay at 0 in the spin-liquid regime. Here, the measured open strings are half of the Wilson loops. By considering the behavior of all types of loops and strings—closed and open, Z and X—we find that there are four regimes (I-IV), corresponding to four phases: (I) Z paramagnet, (II) X paramagnet, (III) topological spin liquid (blue), and (IV) a phase which is consistent with strong decoherence effects (gray). In our analysis, the progression from Regime (III) to (IV) appears to be smooth.

5.7 Experimental Realization in Rydberg Atom Arrays

Motivated by these exciting theoretical results and numerical demonstrations, we now show the use of LED for characterization of topological order in a recent experimental realization of a $\mathbb{Z}_2$-spin liquid on a 219-qubit programmable quantum simulator [143]. In the experiment, qubits are encoded in ground states and $n = 70$ Rydberg states of neutral ^{87}Rb atoms and placed in an array on the links of a kagome lattice (Figure 5.8). This model maps onto a dimer model, where each Rydberg atom can be viewed as a dimer covering the two adjacent vertices of the kagome lattice [160]: the Rydberg blockade interaction between nearby atoms enforces a "dimer constraint" by preventing, with high probability, any vertex from being covered by more than one dimer [136].

This dimer model is predicted to support a $\mathbb{Z}_2$-topologically-ordered state of the resonating valence bond (RVB) type, involving the equally-weighted superposition of all dimer coverings [160,117,127]. In this model, Z-stabilizers are given by (-1) times the product of single-qubit Z-operators on the edges touching a vertex, X-stabilizers are given by the product of off-diagonal operators supported on the triangles bordering a hexagon (see Appendix B.6), and the RVB state forms a fixed-point state. An e (resp., m) anyon arises when a Z (X) stabilizer is violated [137,154,163] ‡.

In the experiment, a topologically-ordered state is prepared by quasi-adiabatically adjusting the detuning Δ and Rabi frequency Ω of a global laser drive [143]. The onset of topological order is observed by studying the expectation values of Wilson loops and open strings [59,23,66,160,143]. A state consistent with $\mathbb{Z}_2$ topological order emerges when using a quasi-adiabatic sweep from initial $\Delta/\Omega < 0$ to a final value of Δ/Ω in the range $3.3 \lesssim \Delta/\Omega \lesssim 4.5$. In practice, several factors make quantitative characterization of such states difficult, as they cause the prepared state to differ from the ideal fixed-point state for the dimer model. In particular, the Rydberg interaction Hamil-

‡The (-1) factor for Z-stabilizers ensures stabilizer expectation values of $+1$, because each vertex is touched by exactly one dimer.

tonian is only an approximation of the parent Hamiltonian of the fixed-point state [§]. Moreover, the finite speed of the detuning ramp and experimental imperfections (e.g., off-resonant scattering, laser phase noise, spontaneous emission events) also modify the experimentally created state. These factors correspond to both coherent and incoherent perturbations, similar to those considered in our toric code simulations. As a result, while topological order can be discerned at modest length-scales, the expectation values of large, bare Wilson loop observables have nearly vanishing signal for almost all final values of Δ/Ω (Figure 5.8b,e).

To circumvent these imperfections, we measure LED loops on the experimentally prepared states. Due to the limited experimental system size, it is not possible to consider loops which strictly satisfy the limit where $\xi \ll d \ll L$, resulting in relatively small expectation values for the LED loop operators. Nonetheless, we clearly observe a range of values of Δ/Ω where both Z- and X-loops are amplified, which corresponds to the spin-liquid interval identified in Ref. [143] (blue shaded region in Figure 5.8). In particular, some of the largest loops within the system acquire non-zero expectation values in this parameter regime. To further confirm our findings in this intermediate system size setting, we examine the behavior of open Z- and X-strings under LED; upon considering the flow of all types of loop and string operators—closed and open, Z and X—we find that there are four regimes (I-IV). Regimes I, II, and III correspond to the Z-paramagnet, X-paramagnet, and spin-liquid regime, in agreement with the experimental paper [143]. We emphasize that our analysis of Regime III goes beyond that of [143], showing non-trivial coherence in closed loops at significantly longer length-scales. In Regime IV, our analysis appears to be consistent with a decoherence-dominated disordered phase, instead of a valence-bond solid (VBS) phase, predicted for the ground state. In particular, we find that the amplification of X-loops continuously decreases, without any

[§]For example, the $1/r^6$ interaction between Rydberg atoms gives rise to long-range tails in the interaction Hamiltonian. These long-range tails also destabilize the spin-liquid ground state, which could cause a first-order phase transition between regions (II) and (IV) in Figure 5.8. Nonetheless, a spin-liquid state can be prepared by using finite ramp speeds, as was done in the experiments [63,42,160].

amplification of open Z-strings. This is analogous to the disordered part of the mixed-state phase diagram (Figure 5.4c), which has a high density of dephasing (Z) errors.

These results demonstrate that our LED method constitutes a promising approach to enhance the detection and characterization of topological order, both theoretically and experimentally. In particular, as experimental developments in quantum simulation and quantum information processing develop over the next several years and decades, techniques such as LED can become indispensable parts of quantum simulation toolboxes for understanding exotic states of entangled quantum matter.

Working hard is important, but there is something else that matters even more: believing in yourself.

J. K. Rowling, Harry Potter and the Order of the Phoenix

6

Hardware-Efficient, Fault-Tolerant Quantum Computation

Throughout human history, the specialization of labor has led to remarkable advances in knowledge and productivity. Analogously, in quantum information, general-purpose quantum error correction protocols can only be as efficient as the society in which every member has the same

duties. Indeed, despite major efforts across different platforms, most general-purpose approaches for error-corrected quantum computation are still out of reach even for the most advanced systems, due to significant overhead in extra qubits and quantum gates. Motivated by these considerations, we propose and analyze the specialized design of fault-tolerant quantum computation protocols tailored to a quantum computer built from arrays of neutral Rydberg atoms, atoms in which one electron is in an extremely excited state.

Inspired by recent experimental advances in quantum control of arrays exceeding 200 atoms, our work provides the first comprehensive study of the relevant error channels in this system and identifies several decay mechanisms that are challenging to address using traditional, general-purpose techniques. We exploit the specific structure of the error model to considerably simplify several error correction requirements, and we make use of important features of neutral atoms to greatly facilitate the key steps in our protocols. These approaches to error correction for neutral Rydberg array quantum computation is dramatically more efficient than existing methods and could be implemented in near-term experiments involving hundreds of programmable atoms. Our results thus open the door towards large-scale fault-tolerant quantum computation using Rydberg atoms.

6.1 Introduction and Motivations

Neutral atom systems have recently emerged as a promising platform for quantum information processing. While the exceptional coherence times of their ground states enable long-lived quantum memories, fast, high-fidelity quantum operations can be achieved by individually addressing atoms with laser pulses and coupling them to highly-excited Rydberg states[84,109,136]. Furthermore, large numbers of individual neutral atoms can be deterministically arranged with arbitrary geometry in two- and three-dimensional systems[7,53,8,9]. Recent experiments have demonstrated quantum manipulation in large arrays of atoms for applications ranging from quantum computing to quantum

simulations and quantum metrology [12,104,94,111,105,65,89,88,124,143,16]. Several latest advances allowing for the dynamic reconfiguration of atoms have even led to realization of logical qubits encoded in color, surface, or toric codes, which is a first step to performing quantum error correction (QEC) on neutral atom platforms [16].

While current experiments are already demonstrating a remarkable level of quantum control, experimental imperfections such as Rydberg state decay will eventually limit the depth of accessible quantum operations. To scale up the computation size, it is therefore essential to consider QEC protocols [122]. In particular, such protocols should be fault-tolerant and protect against the key sources of errors occurring within any of the computation, error detection, and encoding and decoding stages. Multiple fault-tolerant protocols have been proposed for generic quantum platforms [147,92,64,175,38,37,131,35,39], but they do not address certain errors present in Rydberg atom setups. Indeed, Rydberg-atom QEC seems to face a daunting challenge at first glance: Rydberg states could decay into multiple other states, which not only results in leakage errors out of the computational space, but could also give rise to high-weight correlated errors from ensuing undesired blockade effects. Motivated by these considerations, we investigate the effects of these intrinsic errors. Remarkably, by utilizing the unique capabilities of Rydberg systems and the structure of the error model, we can design hardware-efficient, fault-tolerant quantum computation (FTQC) schemes that address these errors despite the aforementioned challenges (Figures 6.1 and 6.2). This tailored FTQC approach can even be much more resource efficient than generic proposals [116,36] (Tables 6.1 and 6.2), which often require a larger number of qubits and quantum operations with smaller threshold error than what is achievable in near-term experiments to perform non-Clifford logical operations, either directly [176,37] or by using state distillation [147,21]. The high overhead associated with such protocols is why experimental demonstrations of QEC have thus far been limited to only one or two logical qubits [123,133,29,47,54].

In this work, we first provide a detailed understanding, from the QEC perspective, of the errors

arising from the finite lifetime of the Rydberg state or from imperfections in Rydberg laser pulses. We then show that nine atoms—seven data qubits and two ancilla qubits—are sufficient to encode each logical qubit fault-tolerantly based on the seven-qubit Steane code [150]; we demonstrate how a universal set of fault-tolerant quantum operations can be performed. For atomic species with sufficiently large nuclear spin and high-fidelity ground-state operations, we show that quantum computation with leading-order fault-tolerance can be achieved even using a simple three-atom repetition code [*]. We find that both the seven-atom and three-atom codes can be implemented on scalable geometries with atoms placed in a triangular lattice configuration (Figure 6.1a,c), allowing for their demonstration and study in near-term experiments.

Our work provides an important advance over prior methods by introducing a novel and distinctly efficient approach to address the leakage of qubits out of the computational subspace. For traditional QEC proposals, such leakage is one of the most difficult and costly types of errors to detect and address, making it unfavorable to encode qubits in large multi-level systems such as neutral atoms. Our method to address these leakage errors makes use of techniques based on optical pumping, such that the multi-level structure of each atom can be utilized as part of the redundancy required for QEC. While we focus on neutral atom-based quantum information processors, these techniques are adaptable to many other hardware platforms—for example, they could also greatly facilitate the correction of leakage-type errors in superconducting qubits or trapped ions. For the Rydberg-atom systems we study, we design a method that even converts all leading-order errors to Pauli-Z type errors (Figure 6.2), which then allows us to develop particularly efficient FTQC protocols.

The chapter is organized as follows: we begin in Section 6.2 by outlining the key insights and

[*]The three-qubit repetition code cannot correct any Pauli-X errors, so it cannot be used for FTQC in typical setups. However, because the error model for Rydberg-atom setups does not contain any Pauli-X errors at the leading order (as shown in Section 6.3), the repetition code is applicable in these platforms. We thus use the term "leading-order fault-tolerance" when describing our Ryd-3 protocol to emphasize this point explicitly.

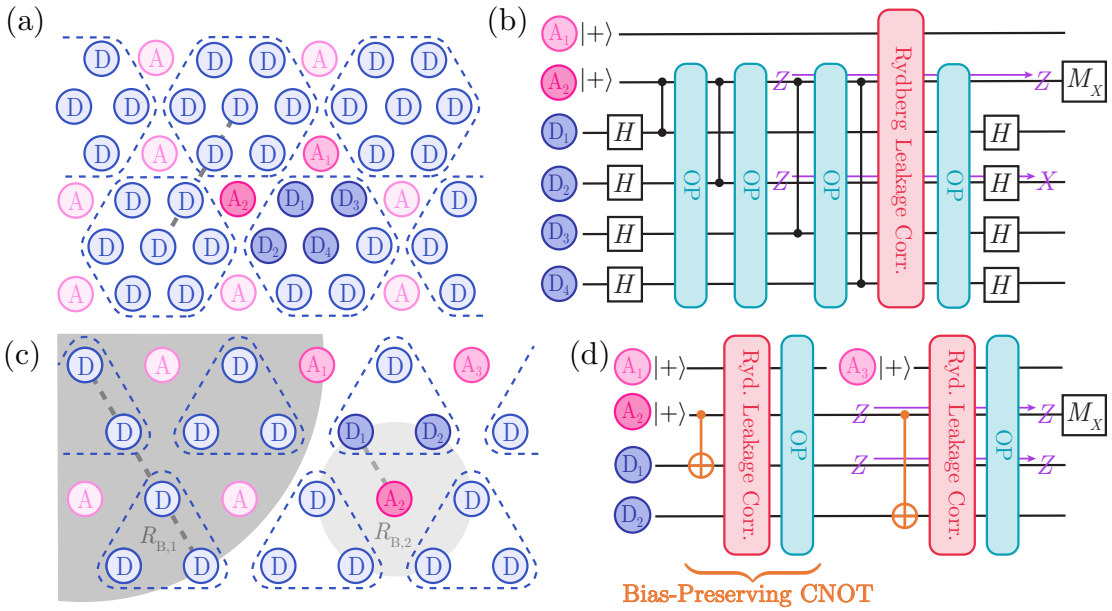


Figure 6.1: Architecture for FTQC with Rydberg atoms. (a) Geometrical layout of atoms for FTQC using the seven-qubit encoding. Data (D, blue) and ancilla (A, pink) atoms are placed on the vertices of a triangular lattice, with seven data atoms comprising a logical qubit (blue dotted hexagons). The dotted grey line indicates the Rydberg interaction range required. (b) Circuit illustrating our procedure to measure a stabilizer operator, $X_1X_2X_3X_4$, for the seven-qubit code supported on the four data atoms highlighted in (a). Optical pumping (light blue, OP) is performed following every controlled-phase gate (black) to correct for leakage into other ground states. Ancilla qubit A_2 (darker pink) measures the stabilizer eigenvalue, while ancilla qubit A_1 (lighter pink) is used to detect and correct for Rydberg leakage errors (red). In this way, all gate errors are converted to Pauli-Z type errors (purple) and do not spread to other qubits. (c) Geometrical layout for quantum computation with leading-order fault-tolerance using the three-atom encoding. Data and ancilla atoms are placed on the vertices of a triangular lattice, with three data atoms comprising a logical qubit (blue dotted triangles). In this case, two Rydberg states with different blockade radii, $R_{B,1}$ and $R_{B,2}$ (dark and light grey, respectively) are required. (d) Our circuit for measuring a stabilizer operator, X_1X_2, of the repetition code, which is supported on the two data atoms highlighted in (c). By combining a novel entangling pulse sequence with Rydberg leakage correction and optical pumping, we implement a bias-preserving CNOT gate (see Figure 6.5), allowing us to perform QEC without introducing X or Y errors at any point in the computation.

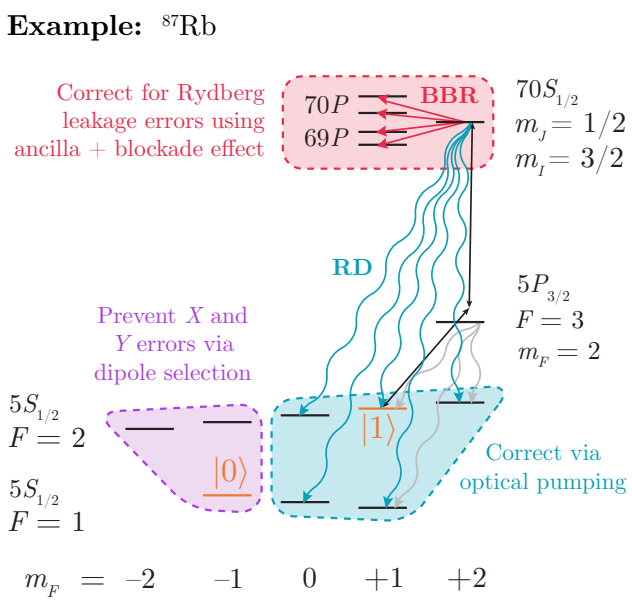


Figure 6.2: Level diagram and dominant error channels form a qubit encoded in the hyperfine clock states of ^{87}Rb. The dominant intrinsic errors for this encoding arise from blackbody radiation (BBR, red), radiative decay (RD, light blue), and intermediate state scattering (grey). Their effects can be determined via dipole selection rules (purple), and the relevant leakage errors can be corrected by making use of the Rydberg blockade effect or optical pumping.

main results of this work. A detailed analysis of the error channels in the Rydberg system is presented in Section 6.3. Under this realistic error model, we design FTQC schemes based on the seven-qubit Steane code in Section 6.4. Furthermore, by utilizing atomic species with high nuclear spin, we develop an alternative, leading-order fault-tolerant protocol in Section 6.5 based on a simple repetition code. We then show in Section 6.6 how the key ingredients of our proposals can be implemented in near-term experiments. Finally, we present conclusions in Section 6.7.

6.2 Overview of Main Results

We consider neutral atoms in a static magnetic field $\mathbf{B} = B_z\hat{\mathbf{z}}$. Due to the nonzero nuclear spin I, the electronic ground state manifold consists of many sub-levels split by hyperfine coupling and a finite $\mathbf{B}$ field. These levels exhibit remarkably long lifetimes, making them particularly good candidates for encoding qubits (or more generally, qudits) for quantum information processing. Furthermore, although neutral atoms in ground electronic states are effectively non-interacting, entangling gates between nearby atoms can be performed by coupling one of the qubit states (e.g. $|1\rangle$) to a Rydberg nS state $|r\rangle$ with large n, which exhibits strong van der Waals interactions (Figure 6.3a). Under certain conditions, these interactions can be interpreted effectively as a blockade constraint prohibiting simultaneous Rydberg population within a blockade radius R_B. These can be leveraged to perform, for example, fast multi-control, multi-target phase gates $R(C_1, C_2, \ldots, C_a; T_1, T_2, \ldots T_b)$ (sometimes also referred to as "collective gates"), which are related to the standard $\mathrm{C}^a\mathrm{Z}^b$ gates upon conjugating all control qubits C_j and the first target qubit T_1 by Pauli-X gates[84,83,105]; this is achieved by applying individually addressed, resonant π and 2π pulses between the qubit $|1\rangle$ state and the Rydberg state (Figure 6.3b). Such an operation is also related to the gate $\mathrm{C}^a\mathrm{NOT}^b$ by single-qubit unitaries and has been demonstrated in recent experiments for small a, b[105].

While this procedure provides an efficient scheme to entangle two or several atoms, for large-

scale quantum computations, the finite lifetime of Rydberg states presents an important source of error even if the rest of the experimental setup is perfect. This lifetime is determined by several contributions. First, interactions with blackbody photons can induce transitions from the nS state to nearby Rydberg $n'P$ states of higher or lower energy; such errors are subsequently referred to as blackbody radiation-induced (BBR) errors. Second, spontaneous emission of an optical frequency photon can result in radiative decay (RD) to a low-lying P state, which will quickly relax into the ground state manifold. In addition, if a multi-photon Rydberg excitation scheme is used for the Rydberg pulses, another intrinsic source of error during Rydberg gates is photon scattering from an intermediate state. These error channels are illustrated in Figure 6.2.

For the purposes of QEC, these errors can be formally described as follows (see Section 6.3): BBR errors give rise to quantum jumps from the qubit $|1\rangle$ state to Rydberg P states (corresponding to a leakage error), as well as Pauli-Z errors within the qubit manifold, while RD and intermediate state scattering may also result in quantum jumps from $|1\rangle$ to the Rydberg nS state or other hyperfine ground states. The relative error probabilities are determined by selection rules and branching ratios. In addition to these intrinsic errors, we also study the errors in the experimental setup such as Rydberg pulse imperfections or finite atomic temperature. We find that these experimental errors fall within a subset of the RD error model and can therefore also be addressed using our techniques. We note that, throughout this work, we assume the rotations within the hyperfine manifold have much higher fidelity than the Rydberg pulses, as is typically the case. Such errors can also be suppressed to high orders by using existing experimental methods such as composite pulse sequences or by incorporating traditional QEC techniques such as concatenation.

6.2.1 Reduction to Pauli-Z errors

To protect against the errors mentioned above, three critical observations are used (see Figure 6.2). First, we note that quantum jumps from $|1\rangle$ to Rydberg states associated with BBR can be detected

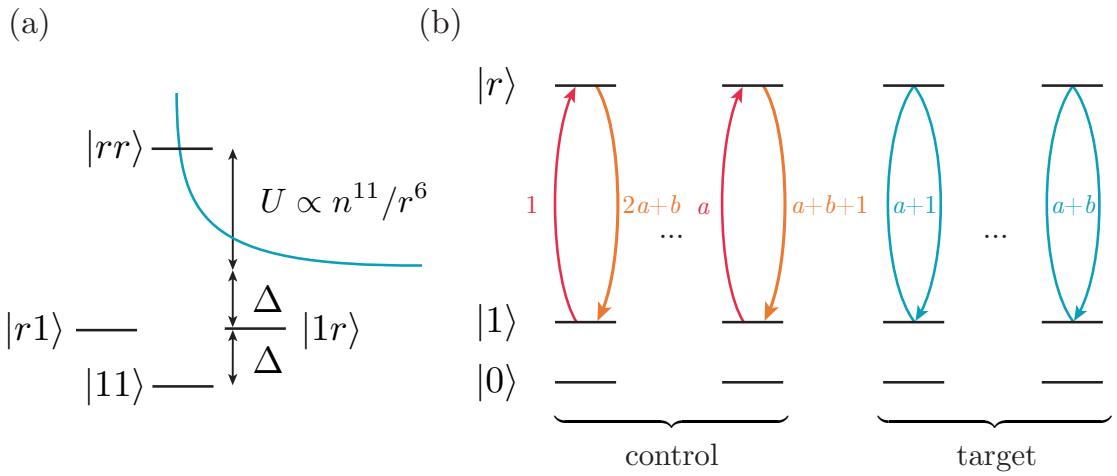


Figure 6.3: Quantum computing with Rydberg atoms. (a) Rydberg blockade mechanism. Δ is the Rydberg laser detuning, and the Rydberg interaction strength $U \propto n^{11}/r^6$, where n is the principal quantum number and r is the atom separation. (b) Protocol for performing a multi-qubit entangling Rydberg gate $R(C_1, C_2, ..., C_a; T_1, T_2, ...T_b)$ on a set of atoms which are all within one given blockade volume. Resonant π pulses $|1\rangle \leftrightarrow |r\rangle$ are first applied to each control qubit (red arrows), followed by 2π pulses on each target qubit (blue arrows). The control qubits are then returned to the ground state manifold via the π pulses shown in orange. Labels on the arrows indicate the ordering of pulses. This Rydberg gate is related to the more conventional controlled-phase gate $\mathrm{C}^a\mathrm{Z}^b$ by conjugating all control qubits and one target qubit by Pauli-X operations, or by applying Pauli-Z gates on both control and target qubits in the special case of $a = b = 1$ ($\mathrm{CZ} = R(C_1; T_1)Z_{C_1}Z_{T_1}$). It can also be used to implement $\mathrm{C}^a\mathrm{NOT}^b$ from $\mathrm{C}^a\mathrm{Z}^b$ by conjugating the target qubits by Hadamard gates. The Rydberg gate $R(C_1, C_2, ..., C_a; T_1, T_2, ...T_b)$ is sometimes referred to as a "collective gate."

via the Rydberg blockade effect by using a nearby ancilla qubit, and subsequently converted to a Pauli-Z type error by ejecting the Rydberg atom and replacing it with a fresh atom prepared in the $|1\rangle$ state[16]. Second, quantum jumps from $|1\rangle$ to ground state sublevels outside the qubit subspace can be corrected via optical pumping techniques. This is particularly efficient as it does not require any qubit measurement for feed-forward corrections, unlike traditional proposals for correcting leakage errors[4]. Third, for atomic species with large enough nuclear spin, dipole selection rules prevent a stretched Rydberg state from decaying to certain ground state sublevels. By making use of this multi-level structure of neutral atoms along with the high-fidelity manipulations of hyperfine states, we can ensure that RD and intermediate state scattering errors do not result in $|1\rangle \rightarrow |0\rangle$ transitions, thereby eliminating X and Y type errors from the error model. This reduction of error

	2-qubit gates	3-qubit gates	Ancillas
7-qubit flagged [38]	36 (48)	0	2
15-qubit flagged [38]	80 (112)	0	2
Ryd-7	**24 (36)**	**0**	**2**
Ryd-3	**8 (16)**	**4 (8)**	**4**

Table 6.1: Comparison of resource costs for leading-order fault-tolerant measurement of all stabilizers. Numbers in parentheses indicate the maximum number of operations required in the unlikely scenario where an error is detected. Details on how to obtain the gate counts for the Ryd-7 and Ryd-3 protocols can be found in Appendix C.8.

	2-qubit gates	3-qubit gates	Ancillas
Yoder, Takagi, and Chuang [176] (CCZ)	162	21	72
Chao and Reichardt [37] (CCZ)	1352 (1416)	84	4
Ryd-7 (CCZ)	**0 (78)**	**27 (29)**	**2**
Ryd-3 (CCZ)	**0 (18)**	**27 (27)**	**4**
Ryd-3 (H)	**20 (28)**	**53 (57)**	**10**

Table 6.2: Comparison of resource costs for the highest-cost fault-tolerant logical operation. CCZ denotes the three-qubit controlled-controlled-phase gate, while H denotes the single-qubit Hadamard gate. Numbers in parentheses indicate the maximum number of operations required in the unlikely scenario where an error is detected. For the Rydberg protocols, the gate counts presented assume a blockade radius of $3d$, where d is the nearest-neighbor lattice spacing. Derivations of the gate counts for the Ryd-7 and Ryd-3 protocols can be found in Appendix C.8, while details on how to obtain the blockade radius requirement can be found in Appendix C.9.

types can significantly alleviate the resource requirement for FTQC.

6.2.2 Fault-tolerant protocols

We now describe two FTQC protocols to address these intrinsic errors in neutral Rydberg atom platforms. The first is based on the seven-qubit Steane code [150], while the second uses the three-qubit repetition code; the latter is more compact and efficient, but has additional experimental requirements such as control over multiple Rydberg states and more complex encoding of logical operations. To realize the seven-qubit code (Ryd-7), we notice that logical state preparation, stabilizer measurements, and a universal set of logical gates (Hadamard and Toffoli [144]) can be implemented using only controlled-phase (CZ) or controlled-controlled-phase (CCZ) gates, up to

single-qubit unitaries at the beginning and end of the operation. For example, while the stabilizer measurements are typically presented as a sequence of CNOT gates between the data atoms and an ancilla atom, these CNOT gates can be constructed by conjugating a CZ gate with Hadamard gates on the target qubit. By mapping each Rydberg gate error to a Pauli-Z error, we therefore ensure that it will commute with all subsequent entangling gates in the logical operation or stabilizer measurement, so it does not spread to other qubits (Figure 6.1b). The resulting single-qubit X or Z error can be corrected by the seven-qubit code in a subsequent round of QEC. This eliminates the need for "flag qubits," which are otherwise necessary to prevent spreading of errors as discussed in Refs. [38,37]. To further reduce resource costs for experimental implementation, we make additional use of the structure of the Rydberg error model, stabilizer measurement circuits, and logical operations of the seven-qubit code. For instance, one of our key findings is that leakage errors into other Rydberg states do not need to be corrected after every Rydberg gate, but can be postponed to the end of a stabilizer measurement (e.g. Figure 6.1b). This allows us to minimize the number of intermediate measurements necessary for each FTQC component, which is typically a limiting factor in state-of-the-art neutral atom experiments.

The simplified error model introduced by conversion of all Rydberg gate errors to Pauli-Z errors motivates us to use the three-qubit repetition code instead of the seven-qubit code to design a leading-order fault-tolerant protocol (Ryd-3). In this case, the stabilizer measurement circuits are also comprised of CNOT gates on data atoms controlled by the ancilla. However, the implementation of each CNOT must be modified: when a CZ gate is conjugated by Hadamard gates as in Figure 6.1b, a Pauli-Z type error that occurs during the CZ gate will be converted to a Pauli-X error after the Hadamard. Such an error can no longer be corrected by the repetition code. Additional errors, such as radiative decay of a control qubit prior to manipulation of the target qubit, can lead to error spreading and correlated errors.

These errors can be addressed via a protocol to directly implement CNOT gates in a *bias-preserving*

way, such that these implementations will not generate any Pauli-X and Y errors to leading order (Figures 6.5 and 6.6). Our protocol makes use of the rich multilevel structure of atoms with large nuclear spin ($I \geq 5/2$, e.g. ^{85}Rb, ^{133}Cs, ^{87}Sr, ...), as well as additional Rydberg states for shelving. Furthermore, we leverage the fact that pulses between certain (i.e. hyperfine) levels can be performed with very high fidelity, so that leading-order errors involve only Rydberg state decay or Rydberg pulse imperfections. This assumption is particularly important, as Ref.[68] shows a no-go theorem stating that a bias-preserving CNOT gate cannot be implemented in any qudit system with a finite number of levels without such structure in the error model. To circumvent this, our pulse sequence directly implements a hyperfine Pauli-X gate on the target qubit only if a nearby Rydberg atom is present (without the need for subsequent Hadamard gates), and we show that errors during this sequence can all be mapped to Pauli-Z errors. Additionally, correlated errors due to control-atom decay can be prevented by using multiple control atoms, such that if one atom decays, the remaining atom(s) still ensure proper gate operation on the target atom. This bias-preserving CNOT protocol can be directly generalized to implement a bias-preserving Toffoli operation, enabling a leading-order fault-tolerant implementation of each operation of the three-atom repetition code. Throughout the chapter, we use the term "leading-order fault-tolerance" in referring to the Ryd-3 protocol as our framework does not inherently address all single-qubit errors, but existing experimental techniques such as composite pulse sequences can be used in conjunction with our protocol to suppress such errors to higher orders (see Section 6.3.4).

Upon comparing our protocols with existing, general-purpose FTQC proposals, we find that the number of required physical qubits and gates for both of our approaches is dramatically reduced (Tables 6.1, 6.2). For example, as seen in Table 6.2, performing the highest-cost operation from our logical gate set, our Ryd-7 protocol requires only 2 ancilla qubits compared with 72 ancillas in Yoder, Takagi, and Chuang[176]. Likewise, Ryd-7 uses at most 60 2-qubit gates (when errors are detected) to perform this logical operation, instead of 1416 gates as in Chao and Reichardt[37]. Such

a significant reduction is possible for our protocols because we leverage both the special structure of the error model and the unique capabilities of Rydberg setups.

Several aspects of the comparison above should be considered. Specifically, we note that certain single-qubit errors addressed in Refs. [176,37] cannot be corrected in our protocols (e.g. Pauli-X errors arising from rotations in the hyperfine manifold). However, we emphasize that Refs. [176,37] also did not consider additional types of errors such as leakage errors which are corrected by our protocol. Indeed, incorporating leakage correction would further increase the resource cost for the earlier proposals considerably. As such, Tables 6.1, 6.2 must be interpreted as a comparison of the cost ensuring fault-tolerance against the leading-order sources of error in a given setup. In the case of Refs. [176,37], these errors include all single-qubit Pauli errors but not leakage errors, while in Rydberg systems, one must address leakage errors at leading order but can neglect certain single-qubit errors.

6.2.3 Towards experimental implementation

For scalable implementation of our FTQC protocols, it is important to consider the geometrical placement of atoms. In addition, because Rydberg entangling gates can only be implemented between atoms within the blockade radius R_B, each protocol defines a minimum value of R_B (in units of d, which is the smallest atom-atom separation). We find that both the Ryd-7 and Ryd-3 protocols can be implemented naturally when the atoms are placed on the vertices of a triangular lattice as shown in Figure 6.1a,c. For both protocols, the required Rydberg gates can be implemented when the blockade radius (R_B for Ryd-7, or the larger radius $R_{B,1}$ for Ryd-3) is greater than $3d$, an interaction range which has already been demonstrated in recent experiments [12]. This requirement can be further reduced in both cases if it is possible to move atoms in between certain operations while preserving the coherence of hyperfine ground states, a capability which has been recently realized [16].

Each component of our FTQC schemes can be implemented in near-term experiments. For neutral alkali-atom systems, recent experiments have already achieved high-fidelity control and entangle-

ment leading to remarkable demonstrations of quantum simulations and computations [104,105,136]. The near-deterministic loading of atoms into lattice structures as shown in Figure 6.1 has already been realized in two and three dimensions [7,53,8,9].

To perform QEC in our protocol, an important requirement is the ability to measure individual qubits and/or detect Rydberg population while preserving coherence in nearby atoms, such that feed-forward correction can be performed. Several approaches for performing fast measurements of individual qubit states in selected atoms can be realized. For example, these selected atoms can be moved into a "readout zone" where their qubit state can be rapidly detected via fast, resonant photon scattering on a cycling transition. Alternatively, one could use arrays with two species (such as two isotopes of the same atom or two different atomic species), where the data atoms are encoded in one atomic species and ancilla atoms are encoded in another species that can be easily measured [179,148]. Finally, the fast detection of Rydberg states has been recently demonstrated in small atomic ensembles using Rydberg electromagnetically induced transparency (EIT) [174]. These could be integrated with the tweezer array platforms currently used for quantum information processing. In these EIT-based procedures, the Rydberg blockade effect translates to clean signatures in the absorption spectrum, and the collectively enhanced Rabi frequency allows for ultrafast detection on a microsecond time scale [174].

While we focus primarily on neutral alkali atoms in this work, significant developments have also been made using alkaline-earth atoms for Rydberg-based quantum computations [110,170]. The clock transition in these atoms allows for high-fidelity qubit encodings, and the large nuclear spin in fermionic species is particularly advantageous for our protocols, so we conclude by discussing how our FTQC schemes can be generalized and applied to these experiments. More detailed experimental considerations are discussed in Section 6.6.

6.3 Error Channels in Rydberg Atoms

In this section, we analyze dominant error mechanisms for quantum operations involving Rydberg a6.2). Because the predominant errors in single-qubit operations can be suppressed to high orders via composite pulse sequences [158,16], we may primarily focus on errors occurring during Rydberg-mediated entangling operations. The decay channels of the Rydberg states include blackbody radiation-induced (BBR) transitions and spontaneous radiative decay (RD) transitions to lower-lying states [14]. Depending on the specific choice of atomic species, another source of error for Rydberg gates can be the scattering from an intermediate state if a two- or multi-photon excitation scheme is used; this is the case for excitation of ^{87}Rb or ^{85}Rb to Rydberg nS states [12]. We will assume these effects are the predominant source of errors that occur during the entangling operations, and we consider contributions to the error model to leading order in the total error probability.

6.3.1 Error modeling for BBR transitions

When a BBR transition occurs on one of the atoms during an entangling gate, it signals that this atom has started in the $|1\rangle$ state, since $|0\rangle$ is not coupled to $|r\rangle$. Such a procedure corresponds to a 'quantum jump' as discussed in, for example, Ref. [31]. The resulting state will predominantly be a nearby Rydberg state $|r'\rangle$ compatible with dipole selection rules. Due to the relatively long lifetimes of Rydberg states, we may assume that the atom will not decay again within the timescale of several Rydberg gate operations, as these would be higher-order processes. In this case, because the states $|r'\rangle$ are not de-excited in the ensuing operations, one serious consequence of BBR quantum jumps is that the remaining Rydberg operations on atoms within the interaction range will be affected by blockade, potentially resulting in multiple, correlated Pauli-Z type errors. Less intuitively, even if a quantum jump does not occur during the gate operation, the atom's state is still modified due to evolution under a non-Hermitian Hamiltonian: it will be more likely that the atom started out in

the $|0\rangle$ state. More details on the theory of quantum jumps can be found in Ref. [31].

For the purposes of QEC, it is useful to express the decay channels in the Kraus operator form, where time evolution of a density operator is given by $\rho \mapsto \sum_\alpha M_\alpha \rho M_\alpha^\dagger$ and the *Kraus operators* M_α satisfy the completeness relation $\sum_\alpha M_\alpha^\dagger M_\alpha = 1$ [129]. For the BBR error model, there will be one Kraus operator

$$M_{r'} \propto |r'\rangle\langle 1| \tag{6.1}$$

for each possible final Rydberg state $|r'\rangle$, where the proportionality constant is determined by the BBR transition rate from $|r\rangle$ to $|r'\rangle$ (see Appendix C.1). In the absence of quantum jumps, the evolution is given by the Kraus map

$$M_0 = \sqrt{1-P}|1\rangle\langle 1| + \sum_{|n\rangle \neq |1\rangle} |n\rangle\langle n|, \tag{6.2}$$

where P is the probability for a BBR transition to occur.

During entangling operations, these BBR errors can give rise to correlated errors. For example, in the Rydberg gates shown in Figure 6.3, a target qubit can only incur a BBR error if the control qubits were all in the $|0\rangle$ state. Thus, for the C^aZ^b gates shown in Figure 6.3, the possible correlated errors may involve one of the Kraus maps $M_{r'}$ or M_0 occurring on one of the qubits, together with Z-type errors on some or all of the remaining qubits involved in that gate.

The rate of BBR transitions from a given Rydberg state nL to another specific state $n'L'$ can be calculated from the Planck distribution of photons at the given temperature T and the Einstein coefficient for the corresponding transition (see Ref. [14]). For ^{87}Rb atoms excited to the $70S$ Rydberg state, there are four dominant final states associated to these BBR errors (see Appendix C.1); these are illustrated in Figure 6.2 as red arrows. The total rate of BBR transitions summed over all possible

final states is[46]

$$\Gamma_{\mathrm{BBR}} = \frac{4k_B T}{3c^3 n_{\mathrm{eff}}^2}, \tag{6.3}$$

where k_B is Boltzmann's constant, c is the speed of light, and n_{eff} is the effective principal quantum number of the Rydberg state which determines its energy[136]: $E_{nL} \propto -1/(2n_{\mathrm{eff}}^2)$. We note that the overall rate of BBR transitions can be suppressed by operating at higher n_{eff} or operating at cryogenic temperatures.

6.3.2 Error modeling for RD transitions

The spontaneous emission events corresponding to RD transitions can be modeled as quantum jumps involving the emission of an optical-wavelength photon. Unlike BBR, however, the resulting state will be a low-lying P state, which will quickly decay back into the ground state manifold. For the stretched Rydberg state of ^{87}Rb, the RD transitions are almost entirely two- or four-photon decay processes to one of the five states in the ground state manifold indicated by light blue arrows in Figure 6.2 (see Appendix C.1 for the precise branching ratios). For the purpose of QEC, we will separately consider the cases of decay into the qubit $|1\rangle$ state and decay into one of the other ground state sub-levels. Because the spontaneous emission event can occur anytime during the Rydberg laser pulse, the first type of decay can result in a final state which is a superposition of $|1\rangle$ and $|r\rangle$. Upon averaging over all possible decay times during the entire pulse (see Appendix C.2), one finds that these errors can be modeled using a combination of Z-type errors and leakage into the $|r\rangle$ state, with the Kraus operators

$$\begin{gathered} \mathcal{M}_0 = |r\rangle\langle r| + \alpha|1\rangle\langle 1| + \beta|0\rangle\langle 0|, \\ \mathcal{M}_r \propto |r\rangle\langle 1|, \quad \mathcal{M}_1 \propto |1\rangle\langle 1|, \quad \mathcal{M}_2 \propto |0\rangle\langle 0|, \end{gathered} \tag{6.4}$$

where α, β, and the proportionality constants depend on the probability for the atom to incur an RD transition to the $|1\rangle$ state and the specific Rydberg pulse being performed.

At the same time, decay to one of the other ground state sublevels shown in Figure 6.2 leads to leakage out of the computational subspace as in the traditional QEC setting (without influencing Rydberg operations on neighboring atoms). That is, for each hyperfine state $|f\rangle \neq |1\rangle$, we have a Kraus operator

$$M_f \propto |f\rangle\langle 1|, \tag{6.5}$$

where the proportionality constant depends on the probability for an RD transition and the branching ratio from $|r\rangle$ to the specific state $|f\rangle$ (see Appendix C.1). Note that due to dipole selection rules, the number of RD channels with non-negligible final state probability is minimized by choosing to couple the $|1\rangle$ state to a so-called "stretched Rydberg state" for entangling gates [†]. In particular, in this analysis, the decay into the qubit $|0\rangle$ state is negligible to leading order. Such an event, corresponding to the Kraus operator $M \propto |0\rangle\langle 1|$ (or equivalently, Pauli-X and Y errors), is considered when we discuss methods to suppress residual errors in our protocols.

As in the BBR case, the absence of quantum jumps results in the atom's population being shifted toward the $|0\rangle$ state, which can be modeled using Pauli-Z errors. RD errors can also give rise to correlated errors when they occur during the primitive entangling gates illustrated in Figure 6.3. In this case, possible correlated errors may involve one of the aforementioned Kraus maps occurring on one of the qubits, together with Pauli-Z and/or $|r\rangle\langle 1|$ errors on some or all of the remaining qubits involved in that gate.

While as noted above, the rate of BBR transitions depends upon the temperature T and n_{eff}, the total RD rate is temperature-independent. Due to reduced overlap between the atomic orbitals, it scales as $\Gamma_0 \sim 1/n_{\text{eff}}^3$ [107]. Comparing this with the scaling for the BBR decay rate, we see that

[†] If, for example, we had instead chosen a Rydberg state with $m_J + m_I = 0$, there would be several additional final states in each case.

while both error rates decrease for larger n, BBR processes dominate for large n, and RD processes dominate for smaller n or very low T.

6.3.3 Errors from intermediate state scattering

When multi-photon excitation is used to couple the $|1\rangle$ state to the Rydberg state, scattering from an intermediate state can give rise to another important intrinsic source of error. By using σ^+-polarized light in the first step of the excitation and choosing the intermediate state to be a $P_{3/2}$ state with the lowest possible n, the intermediate state scattering channels form a subset of the RD channels—they can only result in decay into the qubit $|1\rangle$ state or two other hyperfine ground states, as shown in grey in Figure 6.2 [‡]. Thus, whenever intermediate state scattering is not explicitly mentioned in the following sections, we will assume it has been incorporated with RD errors. We also note that this error rate can be suppressed by increasing intermediate laser detuning in the multi-photon transition, while also increasing laser power.

6.3.4 Experimental imperfections

While BBR, RD, and intermediate state scattering processes constitute the dominant errors for Rydberg-mediated collective gates, it is also important to consider other forms of error, such as technical imperfections in the experimental setup. As discussed in Refs. [12,104,16], the most significant errors of this kind are atom loss and fluctuations in laser phase, intensity, and frequency. The Rydberg laser fluctuations can all be modeled using Pauli-Z errors and leakage into the $|r\rangle$ state, so these errors can be addressed together with the other errors discussed above. Finite atomic temperature, resulting in velocity spread and Doppler broadening on the Rydberg transition [104], likewise leads to

[‡]If an intermediate state with higher n is used, such as $6P_{3/2}$, this is still true to leading order; however, a (highly improbable) four-photon process could potentially lead to mixing between the qubit states $|1\rangle$ and $|0\rangle$.

Pauli-Z errors and leakage into the $|r\rangle$ state. Temperature-induced positional spread causes similar errors, and due to the robustness of the blockade-based gate, these errors can even be rendered negligible with sufficiently large interaction strengths [105]. On the other hand, atom loss forms a more complicated version of a leakage error (called *erasure* in the quantum information literature [64]). However, as discussed in Appendix C.3, we find that such errors can also be addressed efficiently in the present framework. In certain cases, the special properties of these errors can be further leveraged to improve QEC efficiency, as done in the recent proposal of Ref. [173].

Experimental imperfections can also affect the hyperfine qubits used for storing quantum information and performing single-qubit gates; however, these primarily result in Pauli-Z errors and leakage to other hyperfine states, which group together with the error types described above. Moreover, these tend to be significantly smaller sources of error than the two-qubit gates [16]. By choosing a magnetically insensitive transition for our qubit states, we eliminate the leading order errors arising from magnetic field fluctuations. However, Z-type dephasing errors can still arise from the differential light shift from the optical trap. Finite atomic temperature, fluctuating tweezer power, and atom heating can thus cause dephasing, although these can be alleviated to achieve qubit coherence times $T_2 \sim 1$ s by applying standard dynamical decoupling sequences [16]. Leakage to other hyperfine m_F states can also occur due to so-called Raman scattering from the tweezer light, but these effects can be greatly suppressed to timescales $>$ 10 s by sufficiently detuning the tweezer light [16]. Since our qubit states are separated by $\Delta m_F = 2$ (a nuclear-spin-flip transition), bit-flip X and Y error rates from tweezer-induced scattering are even smaller. Finally, temperature-induced Doppler effects, which could in principle result in Z-type errors, are negligible since the qubit transition is of microwave-frequency, and microwave phase stability can be exceptional on the Raman laser used for single-qubit manipulations.

At the same time, as noted earlier, certain experimental imperfections associated with the hyperfine rotations are not directly corrected with our protocol, but can be minimized or suppressed via

other mechanisms such as composite pulse sequences. For example, the primary source of single-qubit gate errors in recent experiments involves laser amplitude drifts or pulse miscalibrations, which can result in X, Y, and Z-type errors [16]. However, these coherent errors can be significantly suppressed by using composite pulse sequences, as done in Ref. [16]: in particular, the BB1 pulse sequence suppresses pulse amplitude errors to sixth order [158]. On the other hand, the error rates associated with phase noise in single-qubit gates are typically much smaller: for example, the phase noise in $^{171}Yb^{+}$ hyperfine qubits has been shown to limit coherence to order 5000 seconds [168]. Although other sources of frequency flucutations result in a T_2^* of approximately 4 ms for the Rb qubit of Ref. [16], thereby inducing pulse frequency errors, these errors are strongly suppressed to second-order due to the MHz-scale Raman Rabi frequencies, and they can be further suppressed with improved cooling and microwave source stability. Furthermore, they can be made completely negligible by using appropriate composite pulse sequences [158]. Finally, incoherent scattering from the Raman beams used for single-qubit rotations can also cause leakage and X, Y-type errors, which can be on the 10^{-5} level [16] for far-detuned Raman beams used for electron-spin-flip transitions but may be higher for nuclear-spin-flip transitions as used for the qubit states here. These remaining hyperfine qubit error rates are significantly smaller than the primary sources of error considered, and they can be further corrected via concatenation of additional error correction codes.

6.3.5 Summary of error channels

We have shown that the multi-level nature of neutral atoms gives rise to various complexities in the error model, including a large number of decay channels and the possibility for Rydberg leakage errors to influence many future operations, resulting in high-weight correlated errors. Despite these complications, one important feature of our error model makes it substantially simpler than the set of all Pauli errors studied in more generic setups—no Pauli-X or Y-type errors are introduced during our Rydberg gates. Indeed, in the following sections, we will show how all the additional

leakage errors and correlated errors in our error model can be converted into Z-type errors, and we use this to design FTQC protocols with substantially reduced resource costs. This reduction to Pauli-Z errors can be found in Sections 6.4.1, 6.4.2 for the seven-qubit code and Section 6.5.1 for the repetition code.

6.4 FTQC with the Seven-Qubit Steane Code

Having established the error model for the Rydberg operations, we now proceed to develop fault-tolerant schemes to detect and correct these errors and perform a universal set of logical operations. The key concept for this construction is the ability to convert all errors described in the previous section into Pauli-Z type errors by introducing ancilla qubits and using the blockade effect, dipole selection rules, and optical pumping (see Figure 6.2). We begin by demonstrating the protocol when only BBR errors are significant (i.e. in the limit of higher Rydberg principal quantum number n), as the error model and QEC mechanisms are simpler to understand in this case. The universal gate set we develop comprises a logical Hadamard gate and a logical controlled-controlled-phase (CCZ) or Toffoli gate [144]. We then describe the more general case involving both BBR and RD errors. Subsequently, we compare the resource cost of our protocol against other fault-tolerant computation schemes and discuss considerations for scalable computation. The final scheme we present in this section is referred to as Ryd-7. Throughout this section, we will use qubits encoded in ^{87}Rb as a concrete example to illustrate our protocols.

While various equivalent definitions of FTQC have been given in the literature for traditional error models, to accommodate the possibility of Rydberg leakage errors—that is, any Rydberg population remaining after the gate operation—we must use the following, stricter one:

Definition 7. A distance-d QEC code is *fault-tolerant* if after any round of error detection and correction, to order $(p_{\text{tot}})^t$, at most t single-qubit Pauli errors are present, where $t = \left\lfloor \frac{d-1}{2} \right\rfloor$ and

p_{tot} is the sum of all error probabilities. In addition, no Rydberg population can be present after any round of error detection and correction.

The final requirement is important because any remnant Rydberg population could blockade future Rydberg gates.

In the following, we will examine the case of code distance $d = 3$ and $p_{\text{tot}} \sim O\left((\Gamma_{\text{BBR}} + \Gamma_0)/\Omega\right)$. Our QEC proposal has the following properties: to leading order in p_{tot},

1. Code states can be prepared with at most a single physical qubit error, without leaving any final Rydberg state population.

2. After each round of error detection and correction, there is at most a single physical qubit error per logical qubit, and there is no Rydberg state population.

3. Each logical gate introduces at most a single physical qubit error per involved logical qubit, without leaving any final Rydberg state population.

It is straightforward to show that any distance-3 code satisfying the above properties is fault-tolerant.

Throughout the rest of the chapter, we will use the term *data qubit* to refer to physical qubits used to encode a logical qubit, and *ancilla qubit* for physical qubits which are used to perform stabilizer measurements or detect errors.

6.4.1 FTQC with BBR errors

Qubit encoding

Our quantum code is based on the seven-qubit Steane code, introduced in Chapter 3, which uses a logical state encoding derived from classical binary Hamming codes [150]:

$$\begin{aligned}|0\rangle_L = &\frac{1}{2\sqrt{2}}(|0000000\rangle + |1010101\rangle + |0110011\rangle + |1100110\rangle \\ &+ |0001111\rangle + |1011010\rangle + |0111100\rangle + |1101001\rangle)\end{aligned} \tag{6.6}$$

$$\begin{aligned}|1\rangle_L = &\frac{1}{2\sqrt{2}}(|1111111\rangle + |0101010\rangle + |1001100\rangle + |0011001\rangle \\ &+ |1110000\rangle + |0100101\rangle + |1000011\rangle + |0010110\rangle).\end{aligned} \tag{6.7}$$

The stabilizer operators for this code are

$$\begin{aligned} g_1 &= IIIXXXX & g_2 &= IXXIIXX & g_3 &= XIXIXIX \\ g_4 &= IIIZZZZ & g_5 &= IZZIIZZ & g_6 &= ZIZIZIZ. \end{aligned} \tag{6.8}$$

In Equation (6.8) and the rest of the chapter when appropriate, we omit tensor product symbols and qubit indices and assume that the j^{th} operator in each product acts on qubit j. Measurements of the stabilizers $g_1, \ldots, g_6$ allow for unique identification and correction of single-qubit X and Z errors. For instance, the absence of any error corresponds to all stabilizers $g_j = +1$, and a Z error on the first qubit would be detected by $g_3 = -1$ and $g_j = +1$ for all $j \neq 3$. The error can then be corrected via an appropriate single-qubit gate.

Error detection and correction

To fault-tolerantly detect and correct for the errors associated with BBR events, we must be able to address both Rydberg leakage and Pauli-Z errors. For the former case, even though leakage errors in traditional QEC settings can be particularly difficult to detect and correct, the particular form of leakage caused by BBR errors make them much easier to identify—we can use an ancilla and the blockade effect to detect the leaked Rydberg population. Specifically, we prepare a nearby ancilla qubit in the state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and apply a 2π Rydberg pulse to detect whether there is another Rydberg atom within the blockade radius. Due to the blockade effect, the ancilla will be in the $|+\rangle$ (respectively, $|-\rangle$) state if nearby Rydberg population is (is not) present.

Once detected, such errors can be easily converted to atom loss errors or Z-type errors. To convert the error to an atom loss error, we notice that the Rydberg atom naturally expels itself due to the anti-trapping potential of the tweezer[16], and can also be directly ejected in $\sim$ 100 ns by pulsing a weak, ionizing electric field ($\sim$ 10 V/cm[45,10,173]) which removes the ion and electron. The exact location of the ejected atom can be determined by following the atom loss protocol outlined in Appendix C.3 and Figure C.2; subsequently, the error can be corrected by replacing the ejected atom with a fresh atom prepared in the $|1\rangle$ state[16] (thereby converting it to a Z-type error) and applying another round of QEC. To reduce the need for applying the atom loss correction protocol, one could add a preventative step after every entangling gate which incoherently re-pumps any remnant population in several most probable Rydberg states into the qubit $|1\rangle$ state. This procedure, along with more details on the conversion of Rydberg population errors, is further discussed in Appendix C.4.

For fault-tolerant error detection and correction, it is important to note that the ancilla used to probe for Rydberg population may also incur a BBR error. This can be resolved by repeating the detection protocol upon finding a BBR error and also using a multi-step measurement procedure

for the ancilla qubit; details are given in Appendix C.5. Such a protocol will be assumed in all future sections when we use an ancilla to detect for Rydberg population.

To fault-tolerantly detect and correct for Pauli errors, we must measure the stabilizers (6.8) in a manner robust against errors that may occur during the detection procedure. The stabilizers for this seven-qubit code are either products of Pauli-X operators or products of Pauli-Z operators, since the Steane code is a CSS code[64]. The traditional (non-fault-tolerant) way to measure a product of four Pauli-X operators (i.e. stabilizers g_1, g_2, or g_3) uses four controlled-phase gates conjugated by Hadamards (Figure 6.1b, black parts). Since Rydberg gate errors can occur during this protocol, we utilize a second ancilla qubit to detect for BBR errors after each entangling operation and convert them to Z-type errors when detected.

The Z errors that occur during a Rydberg gate (or result from conversion of a BBR error) commute with the remaining CZ operations. Thus, the only errors that can occur during a round of stabilizer measurements, to first order in p_{tot}, consist of a Pauli error acting on the ancilla and a Pauli error on one of the data qubits (Figure 6.1b). By resetting the ancilla and repeating the measurement protocol when a -1 measurement outcome is obtained, we can eliminate the effect of the error on the ancilla qubit. An analogous method can be used for the Z stabilizers. In this way, after each round of stabilizer measurements, the correct stabilizer eigenvalues can be obtained to leading order in p_{tot}, while introducing at most one physical qubit X or Z error.

While we have presented the fault-tolerant stabilizer measurement protocol in the simplest form where Rydberg state detection is performed after every physical gate, this is in fact not necessary. Indeed, if we postpone all such detection operations to the end of a circuit which measures the stabilizer $X_\alpha X_\beta X_\gamma X_\delta$ (where Rydberg gates are applied to data atoms in the order $\alpha, \beta, \gamma, \delta$), the only possible correlated errors that can arise are $X_\beta X_\gamma X_\delta$, $X_\gamma X_\delta$, or X_δ, corresponding to BBR transitions on data atoms β, γ, or δ, respectively. For the stabilizers of Eq. (6.8), these errors will all give rise to distinct error syndromes upon measuring $Z^{\otimes 4}$ stabilizers and can thus be corrected (see Appendix

C.6). This can substantially reduce the number of measurements required to implement our protocol, making it more feasible for near-term experiments. A similar procedure can be applied to measure the $Z^{\otimes 4}$ stabilizers.

Logical Operations

Logical Hadamard, Paulis, and S gate. One particular advantage of the Steane code is the transversality of the logical Hadamard, Pauli, and $S = \mathrm{diag}(1, i)$ gates [150]. Specifically, the logical Hadamard simply consists of a Hadamard on each physical qubit:

$$H_L = \otimes_{j=1}^{7} H_j. \tag{6.9}$$

These operations can be performed without ever populating the Rydberg state, and hence without introducing Rydberg gate errors. Similar decompositions exist for the S gate and the Pauli gates X, Y, and Z.

Logical controlled-phase gate. The controlled-phase gate in the Steane code is also transversal [150]:

$$\mathrm{CZ}_{AB} = \bigotimes_{j_A=j_B=1}^{7} \mathrm{CZ}(j_A, j_B). \tag{6.10}$$

We can thus implement a logical controlled-phase operation by performing only seven physical controlled-phase operations and probing for BBR errors in between each physical controlled-phase gate (to convert them to Z-type errors). This eliminates the possibility of correlated multi-qubit errors within a single logical qubit.

Logical Toffoli gate. To implement the Toffoli gate fault-tolerantly and complete our universal gate set, we implement the logical CCZ gate where the target qubit has been conjugated by Hadamard gates. While this gate is not transversal in the Steane code, it may still be decomposed

into a product of physical CCZ gates in a round-robin fashion [176] (see Appendix C.10 for a derivation):

$$\mathrm{CCZ}_{ABC} = \prod_{j_A,k_B,l_C\in\{1,2,3\}} \mathrm{CCZ}(j_A, k_B, l_C), \tag{6.11}$$

so that a logical CCZ operation can be implemented using 27 physical CCZ operations. In the Rydberg setup, this is implemented with the three-qubit Rydberg gate

$$R(j_A, k_B; l_C) = \mathrm{diag}(1, -1, -1, -1, -1, -1, -1, -1) \tag{6.12}$$

and conjugating all involved data qubits by Pauli-X. To avoid propagation of correlated errors resulting from an input X error which does not commute with these Rydberg gates, we begin by fault-tolerantly measuring all the $Z^{\otimes 4}$ stabilizers, and correcting any detected errors; it is simple to verify that this protocol can only result in single-qubit Z errors. This can also be achieved in a more resource-efficient manner by requiring that the stabilizer measurements immediately preceding every logical CCZ gate be done in a way which measures all $Z^{\otimes 4}$ stabilizers last. Furthermore, Rydberg population detection (followed by conversion to Z-type errors, if necessary) is performed after every Rydberg gate, but stabilizers do not need to be measured until the very end; this is because only Z errors occur during the gate operations. In this way, the logical CCZ satisfies the fault-tolerance property.

Although the physical implementation of the CCZ gate is not transversal, the physical gates may be reordered as they all commute with each other. In doing so, we can eliminate some but not all of the intermediate Rydberg population detection steps, to reduce the total number of measurement operations as we did for the fault-tolerant stabilizer measurements. Specifically, we group the three-qubit physical Rydberg gates of the protocol into nine groups of three, $\mathcal{G}_1, \ldots, \mathcal{G}_9$, so that each physical qubit $j_A, k_B, l_C \in \{1, 2, 3\}$ is used in every group. One example of such a grouping $\mathcal{G}_1, \ldots, \mathcal{G}_9$

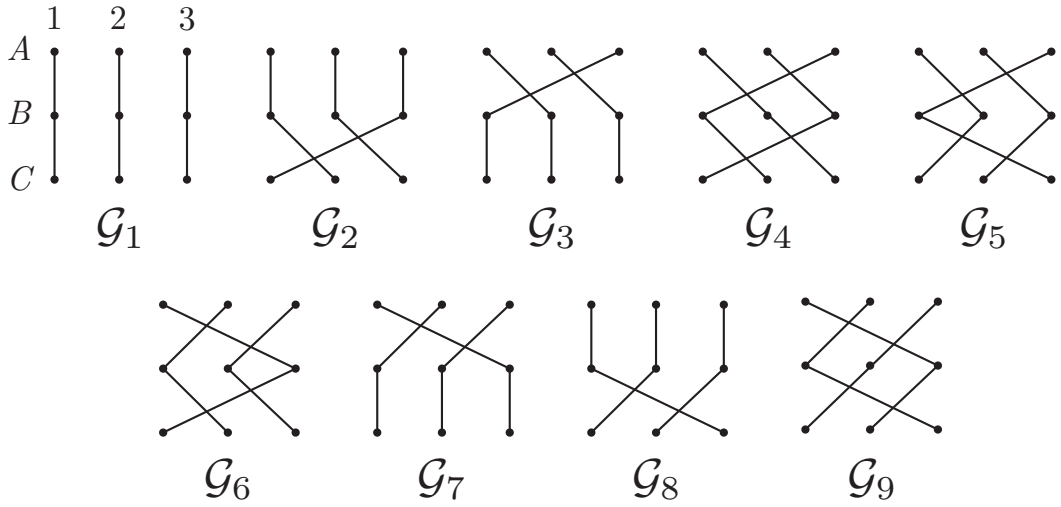


Figure 6.4: Reordering of physical gates in performing the logical CCZ operation. For each logical qubit, only the first three data qubits are shown, since the other data qubits are not involved in the logical gate. Within each group $\mathcal{G}_i$, the Rydberg gates $R(a, b; c)$ are ordered by increasing index of the physical control qubit a (i.e. the data qubit of A involved in the gate).

is shown in Figure 6.4. With this reordering, detection for Rydberg leakage only needs to be performed after each group $\mathcal{G}_i$. This is because a Rydberg leakage error can only result in the blockading of the last two, the last, or no Rydberg gates within a group $\mathcal{G}_i$, and these cases correspond to disjoint possible sets of stabilizer eigenvalues (g_2, g_3) for the three logical qubits (see Appendix C.6).

The Hadamard and CCZ gates together form a universal gate set for quantum computation [144], so we have demonstrated a scheme to construct any quantum operation on the code space fault-tolerantly against BBR errors.

Logical state preparation

Finally, we show that we can prepare the logical $|0\rangle_L$ state in a fault-tolerant manner. The most straightforward preparation of this state uses Steane's Latin rectangle encoding method, whose circuit is shown in Figure 6.5 [150]. In the Rydberg setup, we replace controlled-NOT gates by Rydberg controlled-phase gates with target qubit conjugated by Hadamard gates. Because the Z errors asso-

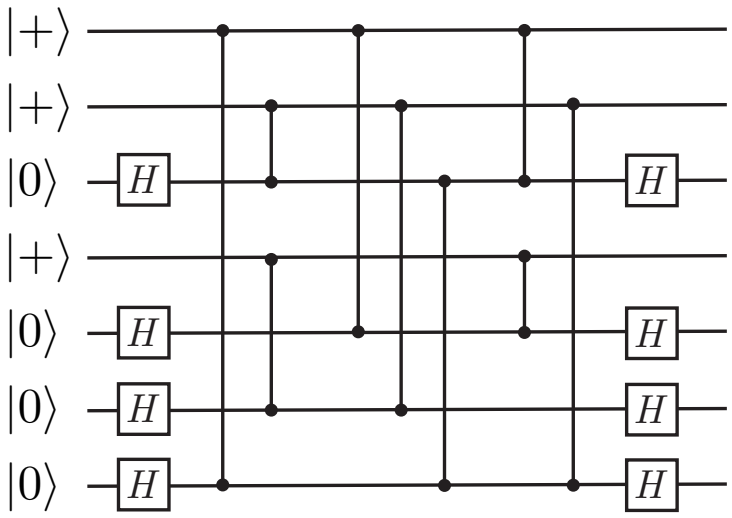


Figure 6.5: Protocol to prepare the logical $|0\rangle_L$ state for the Steane code.

ciated with Rydberg gates commute with controlled-phase operations, to leading order in p_{tot}, there will be at most one Pauli-Z error among the three data qubits initially in the $|+\rangle$ state, and at most one Pauli-X error among the four data qubits initially in the $|0\rangle$ state. Although this could be a two-qubit error, it is correctable because the Steane code identifies and corrects X and Z errors separately. In this procedure, we have assumed we detect for Rydberg population arising from BBR errors after each physical entangling gate and convert these errors to Z errors as necessary. In this way, by applying one round of stabilizer measurements and error correction, we will obtain (to leading order in p_{tot}) a logical $|0\rangle_L$ state with a Pauli error on at most one physical qubit.

6.4.2 FTQC with BBR and RD errors

To address RD errors and intermediate state scattering, we must consider two new classes of leakage errors following the discussion of Section 6.3: (1) leakage into the original Rydberg state $|r\rangle$ and (2) leakage into the other hyperfine ground states, which we will also call "non-Rydberg leakage." The first class of errors is similar to the quantum jumps in the BBR error model, and can be detected and

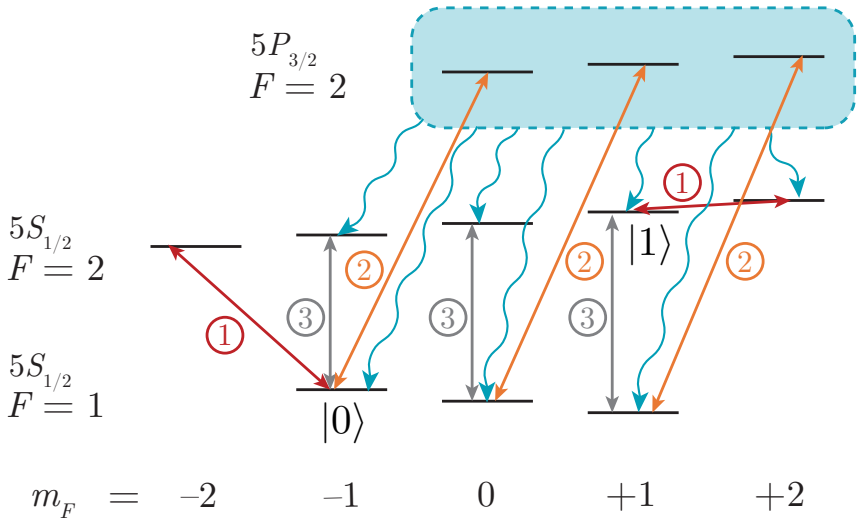


Figure 6.6: Illustration of the optical pumping protocol to convert non-Rydberg leakage errors to Pauli-Z errors in a ^{87}Rb atom. First, we apply π pulses $|1\rangle \leftrightarrow |F = 2, m_F = 2\rangle$ and $|0\rangle \leftrightarrow |F = 2, m_F = -2\rangle$ (red arrows). In the second step, we use σ^+ light to excite states in the $F = 1$ ground state manifold to the $5P_{3/2}$ $F = 2$ manifold (orange arrows). These states decay quickly back into the ground state manifold, as indicated by light blue wavy arrows. Thirdly, we apply resonant π pulses $|F = 2, m_F\rangle \leftrightarrow |F = 1, m_F\rangle$ (grey arrows). The second and third steps are repeated until all population with $m_F \geq -1$ has been transferred to the stretched state $|F = 2, m_F = 2\rangle$. Finally, the first step (red) is repeated to restore the qubit state populations.

corrected in the same way using an ancilla qubit. In the following sections, we will group this error together with BBR errors and refer to them as "Rydberg leakage" errors.

On the other hand, we demonstrate that leakage to other states in the hyperfine manifold can be converted into Pauli-Z type errors using optical pumping. For example, for ^{87}Rb, we design the novel optical pumping protocol shown in Figure 6.6. One crucial property of this optical pumping procedure is that it does not affect the qubit coherence when there is no error. Furthermore, notice that while leakage in traditional QEC settings may be particularly difficult to address, requiring additional entangling gates or ancilla qubits, the particular multi-level structure of neutral atoms allows for efficient correction of these errors. Notably, this optical pumping can be performed without the need for qubit measurement and feed-forward corrections, allowing for efficient implementation in experiments.

Algorithm 1 Fault-tolerant method to measure $X^{\otimes 4}$ stabilizers for Rydberg 7-qubit code.

1) For each $X^{\otimes 4}$ stabilizer $X_\alpha X_\beta X_\gamma X_\delta$:
   a. Initialize ancilla qubit A_2 to $|+\rangle$ state.
   b. Apply gate $Z_j H_j$ to all data qubits $j \in \{\alpha, \beta, \gamma, \delta\}$.
   c. For each $j \in \{\alpha, \beta, \gamma, \delta\}$, apply the Rydberg gate $R(A_1; D_j)$. If $j = \delta$, use ancilla qubit A_1 to detect for Rydberg population as described in Section 6.4.2; if a Rydberg leakage error is detected, convert it to a non-Rydberg leakage error $|F = I + 1/2, m_F = I + 1/2\rangle\langle 1|$. Finally, use the optical pumping technique of Section 6.4.1 to convert any possible non-Rydberg leakage error into a possible single-qubit Z error.
   d. Apply Hadamard gates to all data qubits $j \in \{\alpha, \beta, \gamma, \delta\}$.
   e. Measure A_2 in the X basis.
   f. If A_2 measurement yields -1, break.
2) If any stabilizers are measured to be -1:
   a. Measure all $X^{\otimes 4}$ stabilizers again, this time in the unprotected way and without checking for leakage. There was either already an error in the input, or an error occurred in the initial measurement process. The resulting outcomes will then be the correct stabilizer values to leading order in p_{tot}.

The correction of non-Rydberg leakage errors can be incorporated into the fault-tolerant protocols of the previous section by performing this procedure between the Rydberg entangling gates. Thus, our protocols from the previous section will be fault-tolerant against generic intrinsic Rydberg decay errors. Furthermore, note that when considering this full error model including both BBR and RD events, it is no longer necessary to swap population between the $|1\rangle$ state and the stretched ground state $|F = I + 1/2, m_F = I + 1/2\rangle$ when addressing Rydberg leakage errors (i.e., one can omit Steps 1 and 3 in Appendix C.4); instead, the Rydberg population can be pumped directly to the $|F = I + 1/2, m_F = I + 1/2\rangle$ state, converting it into a non-Rydberg leakage error which is corrected by optical pumping. The full protocols for fault-tolerant stabilizer measurement, the logical controlled-phase gate, and the logical CCZ gate are given in Algorithms 1-3.

While the above discussion has focused on intrinsic RD errors, the non-intrinsic errors of Section 6.3.4 can also be incorporated into our FTQC protocols. Specifically, the errors resulting from Rydberg laser imperfections such as intensity and phase fluctuations only cause Pauli-Z errors and

Algorithm 2 Fault-tolerant logical CZ for Rydberg 7-qubit code.

1) Apply single-qubit Z gates to all physical control and target qubits.
2) For each $j = 1, 2, ..., 7$:
   a. Apply the two-qubit Rydberg gate $R(C_j; T_j)$.
   b. Use ancilla qubit A_1 to detect for Rydberg population as described in Section 6.4.1.
   c. If a Rydberg leakage error is detected, convert it to a non-Rydberg leakage error $|F = I + 1/2, m_F = I + 1/2\rangle\langle 1|$.
3) Use the optical pumping technique of Section 6.4.2 to convert any possible non-Rydberg leakage error into a possible single-qubit Z error.

single-qubit Rydberg leakage errors, so they are already addressed within our current framework. Similarly, atom loss can be detected by using an ancilla qubit and performing a small leakage detection circuit; this is discussed in Appendix C.3. In this case, if a reservoir of atoms is available, we can also convert the atom loss error into a single-qubit Pauli-X or Z error, for instance by replacing the lost atom with a new atom initialized to the $|0\rangle$ state.

6.4.3 Comparison to existing fault-tolerant quantum computing protocols

To demonstrate the significance of our Ryd-7 FTQC protocol and emphasize the importance of considering specific error models when designing QEC approaches, we now compare our model with existing general-purpose FTQC schemes proposed in Refs. [38,37,176]. Specifically, we compare the costs of measuring stabilizers and implementing fault-tolerant logical operations, using as metrics the number of two- and three-qubit entangling operations required for the physical qubits, and the minimum number of ancilla qubits needed. Details on how these numbers can be obtained for the Ryd-7 protocol are provided in Appendix C.8.

Table 6.1 compares the minimum number of two-qubit gates and ancilla qubits required for fault-tolerant stabilizer measurement (and associated error correction) in various QEC proposals. The results for general-purpose FTQC protocols for the 7- and 15-qubit CSS/Hamming codes

Algorithm 3 Fault-tolerant logical CCZ_{ABC} for Rydberg 7-qubit code.

1) Apply X gate to all physical qubits $j_A, k_B, l_C \in \{1, 2, 3\}$.
2) For each group $\mathcal{G}_i$ of physical three-qubit Rydberg gates to apply (where $\mathcal{G}_i$ are ordered as discussed in the main text or Figure 6.4:
   a. Apply gates in $\mathcal{G}_i$.
   b. Use ancilla qubit A_1 to detect for Rydberg population as discussed in Section 6.4.1. If Rydberg leakage is detected:
      i) Convert this leakage error to a possible single-qubit X error.
      ii) Measure stabilizer eigenvalues g_2 and g_3 for each logical qubit in an unprotected way. This is safe because an error already occurred.
      iii) Apply the appropriate correction circuit for the correlated error (since the possible correlated errors all result in disjoint sets of possible syndromes.
      iv) Measure $Z^{\otimes 4}$ stabilizers for all logical qubits in an unprotected way to detect for a possible single-qubit X error induced by step i) above; correct this error if found.
      v) The remaining three-qubit Rydberg gates needed to implement the logical CCZ operation can all be applied in an unprotected way.
   c. Use the optical pumping technique of Section 6.4.2 to convert any possible non-Rydberg leakage error into a possible single-qubit Z error.
3) Apply X gate to all physical qubits $j_A, k_B, l_C \in \{1, 2, 3\}$.

are based on the "flagged syndrome extraction" procedures presented in Refs. [38,175,37]. For each protocol, we separately present the resource cost for cases without any errors and the worst-case cost when an error is present (numbers in parentheses), as the former case is typically much more probable. While the number of ancilla qubits required is the same for all cases, we find that our protocol requires the smallest number of entangling operations in either case even though we must detect for leakage, an additional kind of error not considered in Refs. [38,175,37].

Similarly, Table 6.2 demonstrates this comparison for the fault-tolerant logical CCZ gate, where our improvements are striking. The general-purpose implementation of this non-Clifford gate for three logical qubits in the 7-qubit Steane code is given by Yoder, Takagi, and Chuang [176]; while this implementation requires only a modest number of physical two- and three-qubit gates, it requires a considerable overhead of 72 additional ancilla qubits, making an experimental demonstration very challenging. On the other hand, while Chao and Reichardt's proposal [37] for a fault-tolerant

Toffoli gate using the [[15, 7, 3]] code significantly reduces the ancilla qubit count, the number of physical entangling operations is substantial. Our protocol uses only 2 ancilla qubits compared with 72 required in Yoder, Takagi, and Chuang [176], while using significantly fewer entangling operations (e.g. 60 two-qubit gates) than Chao and Reichardt [37] (1416 two-qubit gates) even in the unlikely scenario where we must correct for an error. We note that while our protocol does use more three-qubit entangling gates than Ref. [176], such gates are nearly as straightforward to implement as two-qubit CZ gates in the Rydberg atom setup (see Section 6.2).

These results clearly demonstrate the advantage of considering a hardware-specific error model and leveraging the unique capabilities of the Rydberg setup when designing FTQC schemes. In particular, even though we must correct for additional errors not considered in traditional settings, we can still dramatically reduce the required number of entangling gates or ancilla qubits.

6.4.4 Scalable implementation

We now discuss some more details regarding the scalable implementation of our protocols, including potential geometrical layouts of physical qubits, resource trade-offs, and residual error rates.

Geometrical considerations. One particular advantage of the Rydberg atom platform is the flexibility in allowing arbitrary geometrical arrangements of atoms. Motivated by recent experimental demonstrations of near-deterministic loading and rearrangement of neutral atoms into regular lattice structures, we propose scalable FTQC architectures in which logical qubits form a coarser lattice on top of the lattice of physical atoms. For the Ryd-7 scheme, one natural layout in a two-dimensional atomic array setup could comprise placing physical atoms on the vertices of a triangular lattice (Figure 6.1a). In this geometry, the hexagonally shaped logical qubits (dark blue dotted hexagons) form a coarser triangular lattice, with ancilla qubits (A, pink) placed on the edges of this coarser lattice to mediate error correction and logical gates. Fault-tolerant universal quantum computation can be performed if nearest-neighbor logical qubits can be entangled; because

physical entangling gates can only be implemented between atoms within a blockade radius R_B, this defines a minimum required value of R_B in terms of the closest atom-atom separation d. Upon examining the physical gates required to implement the logical operations for the seven-qubit code, we find that the requirement in this case is $R_B > 3d$ (dotted grey line), an interaction range which has already been demonstrated in recent experiments [12]. Details on the derivation of this minimal blockade radius can be found in Appendices C.9 and C.10. This requirement on R_B can be further reduced if atoms can be moved in between certain logical operations while preserving coherence between the hyperfine ground states [16].

Resource tradeoffs. For any experiment, resource trade-offs may be made to minimize the total logical error probability. For instance, if the timescale of one round of measurements is much larger than typical gate times (as is the case in certain atomic setups), one may wish to reduce the number of measurement shots required at the expense of performing additional operations. This can be incorporated into our protocol by incoherently driving Rydberg states to the low-lying P state after each entangling gate to convert any possible Rydberg leakage error into the non-Rydberg leakage $|F = I + 1/2, m_F = I + 1/2\rangle\langle 1|$. In this case, ancilla measurements are no longer necessary to detect and correct for Rydberg leakage errors, but this incoherent pumping would be done after every gate, regardless of whether an error had actually occurred. Alternatively, the number of entangling gates can be further reduced at the cost of additional measurements.

Improvements. The FTQC protocol presented in this section relies upon selection rules which impose restrictions on the possible RD error channels. Specifically, as mentioned in Section 6.3, to leading order in the error probability, we ignored the decay channel $|0\rangle\langle 1|$ arising from RD. Given the low branching ratio (determined numerically to be $\sim 10^{-3}$ in ^{87}Rb, see Appendix C.1) from the stretched Rydberg state to $|0\rangle$, this is already a reasonable assumption; however, several approaches can be taken to suppress the probability of such errors even further. First, this probability can be reduced by a factor of roughly 3 or 4 by employing a "shelving" procedure in which population in

the $|0\rangle$ state is swapped with the stretched ground state $|F = -m_F = I + 1/2\rangle$ before and after each entangling gate, due to the lower branching ratio from $|r\rangle$ to this stretched state. To avoid errors arising from near-degenerate Rydberg transitions in this case, one would also transfer population in the $|1\rangle$ state to $|F = I - 1/2, m_F = 1\rangle$ to perform Rydberg excitation in this case, instead of exciting out of the $F = I + 1/2$ manifold. Moreover, by utilizing higher magnetic fields to reduce the branching ratio for RD processes involving large $|\Delta m_F|$, or by using a species with higher nuclear spin (e.g. ^{85}Rb) where the shelving state can be further separated from the stretched Rydberg state, one can suppress the probability of such errors to even higher orders.

6.5 Leading-Order Fault-Tolerance with a Repetition Code

Given that all Rydberg errors can be converted to the Z-type, one may naturally ask whether the full seven-qubit Steane code is even necessary to detect and correct these errors; in particular, one may be tempted to simply use a three-qubit repetition code in the X basis to detect and correct Z-type errors. In such a code, the logical states are

$$\begin{aligned} |+\rangle_L &= |+++\rangle \\ |-\rangle_L &= |---\rangle, \end{aligned} \tag{6.13}$$

and stabilizer operators are

$$g_1 = X_1X_2, \quad g_2 = X_2X_3. \tag{6.14}$$

However, direct application of such a repetition code for FTQC is challenging even with this biased noise model, as one must be able to implement every physical gate in the encoding, decoding, stabilizer measurement, and logical gate procedures without introducing Pauli-X or Y-type errors at any stage—that is, each gate must be implemented in a *bias-preserving* way. This requirement can easily be satisfied for certain physical gates such as the Rydberg controlled-phase or collective gates

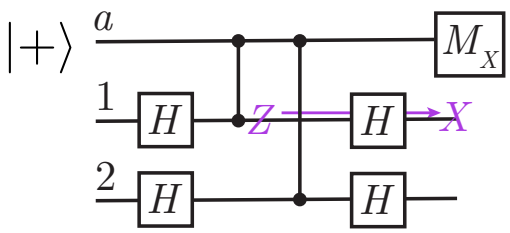


Figure 6.4: Circuit to measure the stabilizer X_1X_2 for the repetition code. CNOT gates must be performed between the ancilla qubit and data qubits 1 and 2. A standard implementation of the CNOT gate using Rydberg controlled-phase gates conjugated by single-qubit Hadamard gates on the target qubits would not be bias-preserving, as a Z error on a target qubit during a controlled-phase gate would become an X error once the final Hadamard gate is applied (purple).

(after all leakage errors are mapped to Pauli-Z type), but is much more difficult to fulfill for other gates. Specifically, measurement of the stabilizers of Eq. (6.14) requires performing controlled-NOT (CNOT) gates as shown in orange in Figure 6.1d.

While a standard implementation of the CNOT gate in a Rydberg setup would comprise the Rydberg controlled-phase gate conjugated by single-qubit Hadamard gates on the target qubit, this would not be bias-preserving: for example, a Z error on a target qubit during a controlled-phase gate would become an X error once the final Hadamard gate is applied (Figure 6.4). In other setups, where a π-rotation of the target qubit about the $\hat{x}$ axis on the Bloch sphere can be performed conditioned on the state of the control qubit (e.g. by engineering a $H_{\text{int}} = ZX$ interaction), an over-rotation or under-rotation error would also translate to an X error and violate the bias-preserving constraint.

Indeed, the implementation of a bias-preserving CNOT may seem unfeasible at first, in light of a no-go theorem proven in Ref.[68]: a bias-preserving CNOT gate is not possible between two qubits encoded in systems where the underlying Hilbert space is finite-dimensional, because the identity gate cannot be smoothly connected to CNOT while staying within the manifold of bias-preserving operations. One way to circumvent this no-go theorem was recently developed for circuit QED systems in Refs.[68,130], where the qubits can be encoded in the continuous phase space of the

photon field, and the dominant source of error—photon loss—can be manipulated via parametric driving schemes to cause only Z-type errors. In our setup, we circumvent the no-go theorem using the special fact that certain pulses in our finite-dimensional atomic system—the pulses between hyperfine states—can be implemented at very high fidelities, so that our leading-order errors arise only from Rydberg pulse imperfections and Rydberg state decay. This allows us to develop a novel laser pulse sequence for entangling Rydberg atoms that directly implements a CNOT or Toffoli gate while preserving the noise bias. Our protocol can be applied on any atomic species with sufficiently high nuclear spin ($I \geq 5/2$). For concreteness, we will illustrate the protocol using the example case of ^{85}Rb throughout the section.

6.5.1 Bias-preserving CNOT in a Rydberg atom setup

As shown in Figure 6.4, the standard implementation of a CNOT gate in a Rydberg system is not bias-preserving. In particular, given the error model for Rydberg gates, X errors on the target qubit can be induced in two ways. First, the target qubit could directly undergo a Rydberg error (e.g. radiative decay) during the controlled-phase gate, resulting in a Pauli-Z error that is transformed into an X error after the Hadamard gate (purple in Figure 6.4). Alternatively, the control atom could decay from the Rydberg state to the ground state at some point during the controlled-phase gate, so that the target qubit Rydberg pulses, which should have been blockaded, are now resonant during the controlled-phase gate. This results in a two-qubit correlated error between the control and target atoms, where the target atom undergoes an X-type error.

Here, we begin by introducing a novel entangling gate pulse sequence for Rydberg atoms to address the target atom X errors. In this discussion, we first assume that the Rydberg pulses on the target atom are either all resonant or all blockaded; that is, we ignore the possibility of a neighboring Rydberg atom decaying during the target atom sequence. We then include this effect and also eliminate the correlated errors by introducing an ancilla qubit and making use of two Rydberg states

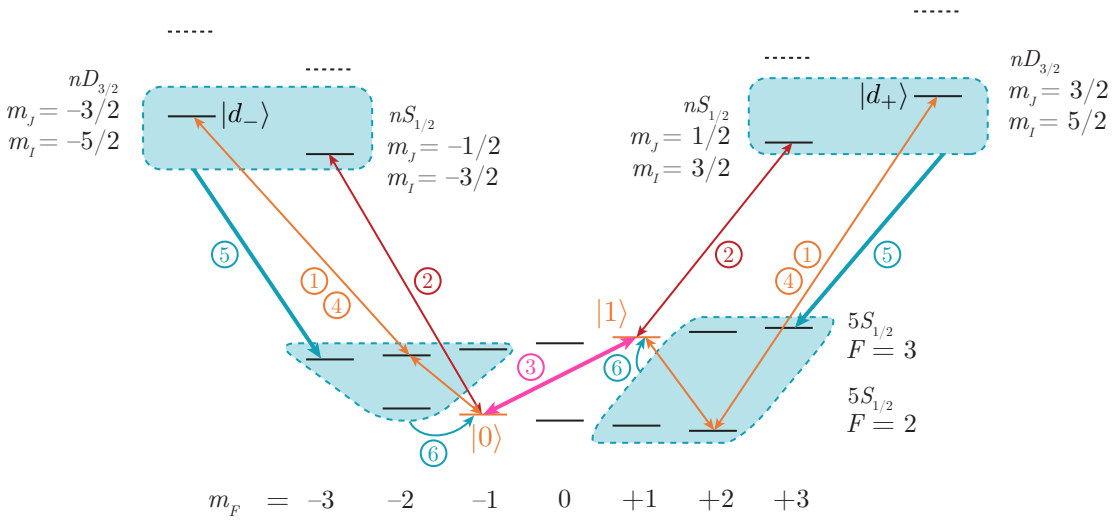


Figure 6.5: Pulse sequence for the target atom in a bias-preserving CNOT gate between ^{85}Rb atoms. Rydberg pulses are resonant if and only if no nearby Rydberg population is present; otherwise, the Rydberg levels are shifted due to the blockade effect (dotted levels). This pulse sequence eliminates target atom X errors in the standard implementation of CNOT shown in Figure 6.4. Step 1: Coherent transfer of population from the qubit states to stretched Rydberg states $|d_\pm\rangle \equiv |nD_{3/2}, m_J = 3/2, m_I = I = 5/2\rangle$. To do this, we first apply hyperfine π pulses $|1\rangle \leftrightarrow |F = 2, m_F = 2\rangle$ and $|0\rangle \leftrightarrow |F = 3, m_F = -2\rangle$, then apply Rydberg π pulses $|F = 2, m_F = 2\rangle \leftrightarrow |d_+\rangle$, $|F = 3, m_F = -2\rangle \leftrightarrow |d_-\rangle$, and finally reapply the hyperfine π pulses $|1\rangle \leftrightarrow |F = 2, m_F = 2\rangle$ and $|0\rangle \leftrightarrow |F = 3, m_F = -2\rangle$ (orange arrows). Step 2: Apply resonant π pulses from the qubit states to the Rydberg states $|1\rangle \leftrightarrow |nS_{1/2}, m_J = 1/2, m_I = 3/2\rangle$ and $|0\rangle \leftrightarrow |nS_{1/2}, m_J = -1/2, m_I = -3/2\rangle$ (red arrows). Step 3: Apply a resonant π pulse between the $|0\rangle$ and $|1\rangle$ ground states (thick pink arrow). Step 4: Repeat Step 1, but use $-\pi$ instead of π pulses on all transitions. Step 5: Incoherently drive any remaining Rydberg population into stretched ground states (thick blue arrows). Specifically, send Rydberg states with $m_J + m_I > 0$ (respectively, < 0) to a stretched $5P$ state with $F = m_F = I + 3/2$ ($F = -m_F = I + 3/2$), which decays quickly and only to the stretched ground state with $F = m_F = I + 1/2$ ($F = -m_F = I + 1/2$). Step 6: Use optical pumping techniques (see Appendix C.12 for details) to map states outside the computational subspace with $m_F > 0$ (respectively, $m_F < 0$) to the qubit state $|1\rangle$ ($|0\rangle$) (thin blue arrows).

with different blockade radii.

To remove the target atom X errors, we wish to design an entangling gate protocol which uses Rydberg states to conditionally swap $|0\rangle$ and $|1\rangle$ population directly, without the change-of-basis from Hadamard gates. This can be accomplished for atomic species with high enough nuclear spin ($I \geq 5/2$). We consider qubits encoded in the ^{85}Rb clock states $|1\rangle \equiv |F = I + 1/2, m_F = +1\rangle$, $|0\rangle \equiv |F = I - 1/2, m_F = -1\rangle$ (orange levels in Figure 6.5), which have a magnetic field-insensitive transition frequency at low fields. The protocol then proceeds as illustrated in Figure 6.5.

The first step of the procedure aims to transfer population in the qubit state $|1\rangle$ (respectively, $|0\rangle$ to the Rydberg state $|d_+\rangle \equiv |nD_{3/2}, m_J = 3/2, m_I = I = 5/2\rangle$ (respectively, $|d_-\rangle \equiv |nD_{-3/2}, m_J = -1/2, m_I = -I = -5/2\rangle$) conditionally, dependent on the state of a control atom. This is achieved because the Rydberg pulses from the qubit states to $|d_\pm\rangle$ are resonant if and only if there are no neighboring atoms in $|r_\pm\rangle$ or nearby Rydberg states. Since each stretched Rydberg state predominantly decays only into ground states with $|\Delta m_F| = |\Delta(m_J + m_I)| \leq 2$ during RD processes, the $|0\rangle$ and $|1\rangle$ populations will not be mixed by Rydberg state decay; however, due to the possible decay channels $|F = 2, m_F = 2\rangle\langle d_+|$ and $|F = 3, m_F = -2\rangle\langle d_-|$, it is possible that the first step fails to excite the atom into a Rydberg state even in the absence of nearby Rydberg population. Consequently, in the second step, we again attempt to transfer the qubit states to Rydberg states, this time using resonant π pulses $|1\rangle \leftrightarrow |nS_{1/2}, m_J = 1/2, m_I = 3/2\rangle$ and $|0\rangle \leftrightarrow |nS_{1/2}, m_J = -1/2, m_I = -3/2\rangle$. Then, in the third step, the population in the qubit states is swapped via the π pulse $|0\rangle \leftrightarrow |1\rangle$. We note that this only swaps population if nearby Rydberg atoms prevented transfer out of the qubit manifold in Steps 1 and 2. Step 4 then acts to invert the first step.

After Step 4, we find that if no Rydberg errors have occurred, the atomic state is restored to the original qubit state (identity map) when no nearby Rydberg population is present, or to the opposite qubit state $|0\rangle \leftrightarrow |1\rangle$ otherwise. Rydberg errors can occur only if the pulses of Step 1 are

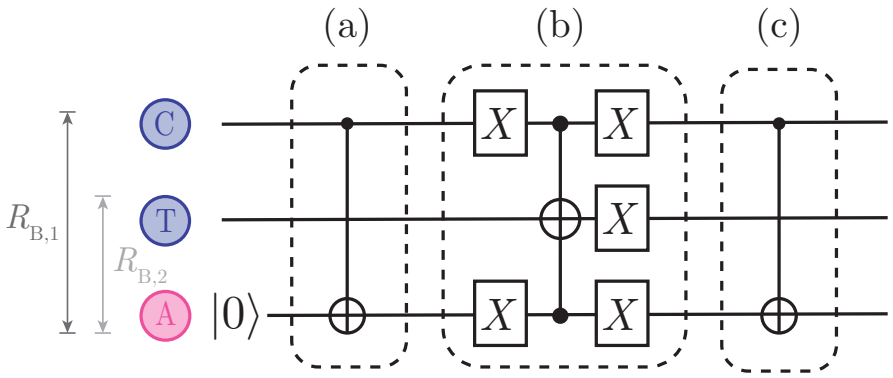


Figure 6.6: Using an ancilla qubit and multiple Rydberg states to eliminate X-type errors arising from control qubit decay. The atoms are positioned on a line, such that atom T is in the middle, and the distance between neighboring atoms is $d \equiv d_{CT} = d_{AT}$. The ancilla qubit is initially prepared in the $|0\rangle$ state. The protocol consists of three steps, labelled (a)-(c), and can be visualized as a quantum circuit. We use two different pairs of Rydberg S states, $|r_{1,\pm}\rangle$ and $|r_{2,\pm}\rangle$, with blockade radii $R_{B,1}$ and $R_{B,2}$, respectively, such that $R_{B,1} > 2d$ and $d < R_{B,2} < 2d$. Steps (a), (c): Apply a CNOT gate with C as control and A as target. This is done by applying a π pulse $|1\rangle \leftrightarrow |r_{1,+}\rangle$ on atom C, performing the pulse sequence of Figure 6.5 on atom A, and applying a $-\pi$ pulse $|1\rangle \leftrightarrow |r_{1,+}\rangle$ on atom C, so that the Rydberg pulses on A are resonant only if C is not in $|r_{1,+}\rangle$ (or a nearby Rydberg state). For these steps, the Rydberg levels $|r_{\pm}\rangle$ in Figure 6.5 are chosen to be $|r_{1,\pm}\rangle$ (see Table 6.3). Step (b): Apply a three-atom gate between C, A, and T. This is done by applying π pulses $|1\rangle \leftrightarrow |r_{2,+}\rangle$ on both atom C and atom A, performing the pulse sequence of Figure 6.5 on atom T, and applying $-\pi$ pulses $|1\rangle \leftrightarrow |r_{2,+}\rangle$ on both atom C and atom A, so that the Rydberg pulses on T are resonant only if neither C nor A is in $|r_{2,+}\rangle$ (or a nearby Rydberg state). For this step, the Rydberg levels $|r_{\pm}\rangle$ in Figure 6.5 are chosen to be $|r_{2,\pm}\rangle$ (see Table 6.3).

resonant (i.e. if no nearby Rydberg atoms are present); moreover, because transitions from $|d_+\rangle$ (respectively, $|d_-\rangle$) only result in states with $m_F > 0$ ($m_F < 0$), any Pauli errors must be of Z-type (for example, projectors $|0\rangle\langle 0|$, $|1\rangle\langle 1|$), and any leakage error must be of the form $|m_F > 0\rangle\langle 1|$ or $|m_F < 0\rangle\langle 0|$. One can then verify that after the pumping steps (5 and 6), the resulting state is the same as in the error-free case, up to a local error of Z type (e.g. $|0\rangle\langle 0|$, $|1\rangle\langle 1|$). As before, the error channels for intermediate state scattering and other Rydberg pulse imperfections can be captured by our error model which contains BBR and RD errors.

Having eliminated X errors arising from target qubit Rydberg errors, we now proceed to address the second type of potential X error arising from control qubit decay. The crux here is to utilize multiple Rydberg atoms (e.g. a control atom and an ancilla atom) to blockade the target atom if the

Step	Rydberg Transitions Addressed		
	Atom C	**Atom T**	**Atom A**
(a), (c)	$\|1\rangle \leftrightarrow \|r_{1,+}\rangle$	none	$\|0\rangle \leftrightarrow \|r_{1,-}\rangle$ $\|1\rangle \leftrightarrow \|r_{1,+}\rangle$
(b)	$\|1\rangle \leftrightarrow \|r_{2,+}\rangle$	$\|0\rangle \leftrightarrow \|r_{2,-}\rangle$ $\|1\rangle \leftrightarrow \|r_{2,+}\rangle$	$\|1\rangle \leftrightarrow \|r_{2,+}\rangle$

Table 6.3: Rydberg transitions used to implement the bias-preserving CNOT gate between two atoms C and T as shown in Figure 6.6. Within each step, one Rydberg transition ($|1\rangle \leftrightarrow |r_{1,+}\rangle$ or $|1\rangle \leftrightarrow |r_{2,+}\rangle$) is addressed for each "control" atom, while two Rydberg transitions ($|0\rangle \leftrightarrow |r_{1,-}\rangle$, $|1\rangle \leftrightarrow |r_{1,+}\rangle$ or $|0\rangle \leftrightarrow |r_{2,-}\rangle$, $|1\rangle \leftrightarrow |r_{2,+}\rangle$) are addressed for each "target" atom. $|r_{1,\pm}\rangle$ and $|r_{2,\pm}\rangle$ have different blockade radii $R_{B,1}$ and $R_{B,2}$ as explained in the main text and in the caption of Figure 6.6.

control is in the $|1\rangle$ state; in this way, if one of the atoms decays, the remaining Rydberg atom(s) can still ensure (to leading order in the total error probability) that the Rydberg pulses on the target atom do not become resonant. For the simplest case, the bias-preserving CNOT gate can be implemented with one ancilla qubit. Let us assume that the control (C), target (T), and ancilla (A) atoms are placed evenly along a line, with the target atom in between the control and ancilla atoms; the ancilla atom is initialized in the state $|0\rangle$. We can make use of two sets of Rydberg states, $|r_{1,\pm}\rangle$ and $|r_{2,\pm}\rangle$, with blockade radii $R_{B,1}$ and $R_{B,2}$, respectively, such that $R_{B,1} > 2d$ and $d < R_{B,2} < 2d$, where d is the distance between neighboring atoms (i.e. between C and T or T and A); as such, atoms C and A are within the blockade radius $R_{B,1}$, but beyond $R_{B,2}$, whereas neighboring atoms are within the blockade radius $R_{B,2}$. The full bias-preserving CNOT gate between the control and target atoms then consists of the three-step procedure illustrated in Figure 6.6, followed by correction of Rydberg leakage errors (as discussed in Appendix C.4) and optical pumping to eliminate non-Rydberg leakage errors (Figure 6.6). The Rydberg transitions addressed in each step of Figure 6.6 are listed in Table 6.3.

This protocol is robust against control atom decay errors, as the Rydberg pulses on atom T are resonant only if neither C nor A is excited to the Rydberg state, and one can see that, to leading order in the total error probability, this can only occur if C starts in the $|0\rangle$ state: first, if C begins in

the $|0\rangle$ state, A must also remain in $|0\rangle$, so the state of T will not be flipped. On the other hand, if C begins in the $|1\rangle$ state and no decay events occur during Step (a), $|C, A\rangle = |1, 1\rangle$ after this step. The Rydberg pulses for T are blockaded in Step (b), so its state will be flipped. Finally, if C begins in the $|1\rangle$ state but decays during the first step, $|C, A\rangle = |1, 1\rangle$ or $|1, 0\rangle$ after this step. The Rydberg pulses for T are still blockaded in Step (b), so its state will be flipped. Finally, Rydberg decay errors in Step (c) will result in projections of the form $|0\rangle\langle 0|$ or $|1\rangle\langle 1|$, which can be expressed in terms of Z errors. In this way, we have eliminated any possible source of X errors arising from the CNOT gate, to leading order in the total error probability. The protocol can also be generalized to implement a bias-preserving Toffoli gate (see Appendix C.7 and Figure C.3). Potential improvements leading to suppression at higher orders are discussed in Section 6.5.3.

The ability to couple atoms to two sets of Rydberg states $|r_{1,\pm}\rangle$ and $|r_{2,\pm}\rangle$ in our bias-preserving CNOT implementation allows atom C to interact with atom A during Steps (a) and (c) of Figure 6.6, but not during Step (b). Alternatively, this tunability of interaction could be achieved with only a single set of addressable Rydberg states $|r_{1,\pm}\rangle$ if the atoms can be rearranged while preserving coherence between hyperfine ground states [16]. In this case, one could move atoms in between Steps (a) and (b) to further separate C, T, and A from each other such that the distance between C and A becomes greater than $R_{B,1}$, while the distance between either of them and atom T remains less than $R_{B,1}$. The atoms can then be returned to their original configuration after Step (b) to allow for interaction between C and A during Step (c).

6.5.2 Leading-order fault-tolerance with the repetition code

The bias-preserving operations discussed above allow for a direct implementation of each component of the three-atom repetition code to perform quantum computation with leading-order fault-tolerance on a Rydberg setup. In particular, logical states (6.13) can be prepared or measured fault-tolerantly in the X basis by transversally preparing or measuring each atom. The measurement

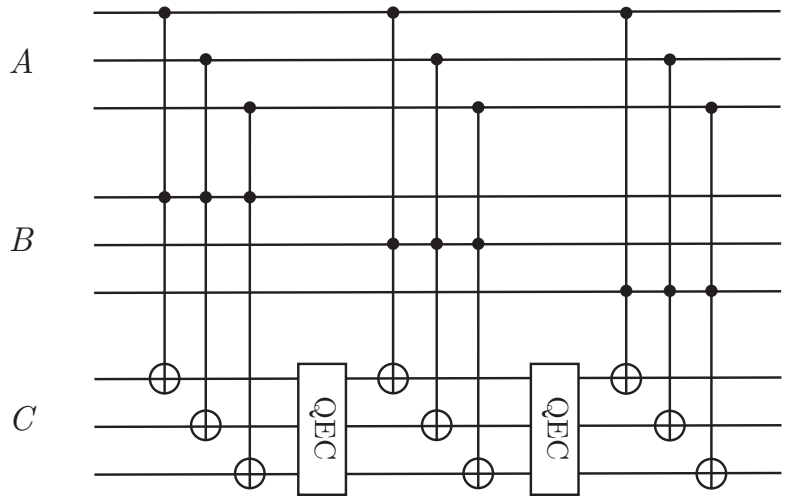


Figure 6.7: Pieceable fault-tolerant implementation of the Toffoli gate in the repetition code.

of stabilizers (6.14) can be achieved using the circuit of Figure 6.1d, where each controlled-NOT gate is done in the bias-preserving way described above; for robustness against errors occurring during this circuit, one must repeat the stabilizer measurement if either g_1 or g_2 is measured to be -1.

A universal set of logical operations can be achieved by implementing a logical Toffoli gate and a logical Hadamard gate as in the seven-qubit case, using the bias-preserving pulse sequences presented above. While not strictly necessary, we will also discuss the implementation of logical controlled-phase and CCZ gates. These may be of use for simplifying the implementation of certain quantum algorithms, as they do not require the new bias-preserving pulse sequences and can be implemented using the standard method for performing Rydberg-mediated entangling gates as described in Figure 6.3.

Logical Toffoli gate. One important feature of the encoding (6.13) is that the logical $|0\rangle_L$ (respectively, $|1\rangle_L$) state consists of an equal superposition of states with an even (odd) number of

physical qubits in the $|1\rangle$ state:

$$\begin{aligned} |0\rangle_L &= \frac{1}{2}\left(|000\rangle + |110\rangle + |101\rangle + |011\rangle\right) \\ |1\rangle_L &= \frac{1}{2}\left(|111\rangle + |001\rangle + |010\rangle + |100\rangle\right). \end{aligned} \tag{6.15}$$

From this observation, one can see that the Toffoli gate CCX_{ABC} with logical control qubits A, B and logical target qubit C can be implemented as a product of nine physical Toffoli gates:

$$\mathrm{CCX}_{ABC} = \prod_{\substack{j_A,k_B\in\{1,2,3\} \\ l_C=j_A}} \mathrm{CCX}(j_A, k_B, l_C). \tag{6.16}$$

Each physical Toffoli gate can be implemented in a bias-preserving fashion as described previously, resulting in at most one physical Z error in each logical qubit, assuming that Rydberg and non-Rydberg leakage errors are converted to possible Z errors after each physical gate. In this case, however, while Z errors on the control qubits A or B would commute with remaining Toffoli gates, a Z error on one of the physical qubits of C could spread to multiple Z errors within A or B after subsequent Toffoli gates if uncorrected. To address this, we order the physical gates as shown in Figure 6.7 and perform error correction after every three physical Toffoli operations by measuring the stabilizers (6.14); this follows the *pieceable fault-tolerant* implementations of non-transversal gates discussed in Refs. [176,68]. In this way, after the entire logical gate, there will be at most one physical qubit Z error per involved logical qubit.

Logical Hadamard gate. Unlike the Steane code, the repetition code is not a CSS code, and its logical Hadamard gate is not transversal. However, as discussed in Ref. [68], the logical Hadamard gate can be implemented using a logical Toffoli gate combined with fault-tolerant measurements in the X basis, as shown in Figure 6.8. The logical Hadamard gate combined with the logical Toffoli or CCZ gate form a universal set of logical operations.

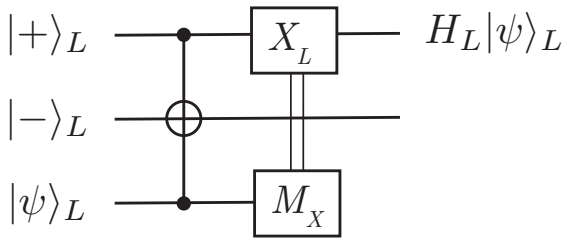


Figure 6.8: Implementing the logical Hadamard in the repetition code using the logical Toffoli gate, as discussed in Ref. [68].

Logical controlled-phase gate. A logical controlled-phase operation in the three-qubit code can be implemented using the standard Rydberg pulse sequences for controlled-phase gates between each pair (j_A, k_B) of physical qubits, where j_A and k_B belong to the encoding of logical qubits A and B, respectively:

$$\mathrm{CZ}_{AB} = \prod_{j_A,k_B\in\{1,2,3\}} \mathrm{CZ}(j_A, k_B). \tag{6.17}$$

To correct for the errors that occur during gates, one should remove any Rydberg population and apply the optical pumping scheme to convert non-Rydberg leakage errors into possible Z errors after each physical controlled-phase operation. The stabilizers only need to be measured after the entire logical operation, since Rydberg gates can only produce Z errors which commute with all the physical CZ gates being performed (and hence do not spread to higher-weight errors).

Logical CCZ gate. Similarly, a logical controlled-controlled-Z operation between logical qubits A, B, C,

$$\mathrm{CCZ}_{ABC} = 1_A 1_B - \frac{1}{4}(Z_A - 1_A)(Z_B - 1_B)(Z_C - 1_C), \tag{6.18}$$

can be implemented as a sequence of physical CCZ operations:

$$\mathrm{CCZ}_{ABC} = \prod_{j_A,k_B,l_C\in\{1,2,3\}} \mathrm{CCZ}(j_A, k_B, l_C). \tag{6.19}$$

As with the case of logical CZ, Rydberg and non-Rydberg leakage errors should be converted to possible Z errors after each physical gate. Notice that even though the logical CCZ is not transversal, this implementation is leading-order fault-tolerant because any given physical gate can result in at most one physical qubit Z error per logical qubit; since Z errors commute with the remaining gates applied, they do not propagate to become multi-qubit errors. While the CCZ gate is not strictly needed for the universal gate set given a leading-order fault-tolerant implementation of the logical Toffoli gate, it requires fewer resources to implement than the logical Toffoli as it uses the standard, simpler Rydberg gates $R(C_1, C_2; T)$ instead of the more complicated bias-preserving CNOT pulse sequences (see Table 6.2). Thus, this operation may be useful for reducing the resource cost of certain quantum algorithms.

6.5.3 Scalable implementation

We now discuss some important considerations for the scalable implementation of our Ryd-3 protocol, including the geometrical layout, resource requirements, and potential improvements.

Geometrical layout. Based on the implementations of logical gates, stabilizer measurement, and the underlying bias-preserving CNOT given in the previous sections, we find that a convenient geometry is to place data and ancilla atoms on the vertices of a triangular lattice as shown in Figure 6.1c, with three data atoms comprising a logical qubit. In this configuration, the logical qubits form a coarser triangular lattice, as in the case of Ryd-7. As discussed in Section 6.5.1, two Rydberg states with different blockade radii $R_{B,1} > R_{B,2}$ are required to implement the bias-preserving CNOT gate. Based on the interaction ranges required for performing fault-tolerant stabilizer measurements and logical operations as described previously, we find that the larger blockade radius must be greater than $3d$ (dark grey in Figure 6.1c, where d is the nearest-neighbor spacing on the square lattice; this is required for some of the physical gates in the logical CCZ and Toffoli gates. On the other hand, the smaller blockade radius $R_{B,2}$ should be strictly between d and $2d$ for ef-

ficient implementation of the bias-preserving CNOT and fault-tolerant stabilizer measurements (light grey in Figure 6.1c). Details on how to obtain the requirement $R_{B,1} > 3d$ can be found in Appendix C.9.

Alternatively, the data and ancilla atoms can be placed on the vertices of a square lattice in an alternating fashion (see Appendix C.11). In this case, the blockade radius requirements are $R_{B,1} > 3.61d$ and $d < R_{B,2} < 2d$. For both the triangular lattice and square lattice geometries, experimental developments allowing for rearrangement of atoms while preserving the coherence of hyperfine ground states could be used to further reduce the requirement on $R_{B,1}$ and eliminate the need for a second set of Rydberg states with blockade radius $R_{B,2}$.

Improvements. While our bias-preserving CNOT suppresses X-type errors to leading order, the amount of bias preservation is ultimately limited by the decay rate of the stretched Rydberg D state into the qubit states. To further suppress these errors, one can shelve to stretched Rydberg states with higher angular momentum, which would have lower decay rate to the qubit states. Alternatively, one can also use an atomic species with higher nuclear spin, where the qubit states can be separated from the stretched Rydberg state by a larger amount $|\Delta m_F|$. Likewise, one could also increase the magnetic field in the experimental setup to suppress the rate of transitions with high $|\Delta m_F|$. To achieve suppression beyond the leading order, one can then use more Rydberg "shelving" states in the target atom pulse sequence of Figure 6.5 and more ancillas to suppress the effects of control atom decay.

The Ryd-3 hardware-tailored FTQC approach inherently addresses errors due to Rydberg pulse imperfections in addition to those arising from the finite Rydberg state lifetime, as these errors fall within a subset of the radiative decay errors. As in the Ryd-7 case, the Ryd-3 approach can also be enhanced to further protect against atom loss errors at the expense of additional physical operations by incorporating the atom loss detection scheme described in Appendix C.3 in between Rydberg operations.

6.6 Further Considerations Towards Experimental Implementation

In this section, we discuss further considerations on how our FTQC protocols can be implemented in near-term experiments. Recent experiments using neutral alkali atom systems have already achieved near-deterministic trapping, loading, and rearrangement of tens to hundreds of atoms into two-dimensional lattice structures such as the triangular lattice needed for our protocol [17,53,7,89]. Furthermore, high-fidelity manipulations within the ground state manifold and two- and three-atom Rydberg blockade-mediated entangling gates have been demonstrated [104,105,65]. Blockade interactions between Rydberg atoms separated by three times the lattice spacing, which is the interaction range required for both of our protocols, have also been realized [12].

6.6.1 Measurements and feed-forward corrections

To perform QEC, an important ingredient is the ability to measure the states of ancilla qubits and/or detect Rydberg population and perform feed-forward corrections. Several approaches can be considered. One promising way to rapidly measure individual qubit states is to resonantly drive a cycling transition and detect the scattered photons [11]. At lattice spacings of a few microns, this detection scheme could be limited by atom heating and cross-talk from the reabsorption of scattered photons by neighboring atoms [151]. To this end, recent developments in coherent transport of entangled atom arrays [16] can be used to mitigate these effects by moving the selected ancillary atom(s) into a detection zone far away from the rest of the array before it is measured.

To estimate the maximum speeds of coherent transport before atom loss and heating become significant, one can consider the harmonic oscillator potential (i.e. the optical tweezer) that the atom is trapped in. Following the analysis in Refs. [33,16], the average energy increase to the atom will be $\Delta E = m|\tilde{a}(\omega_0)|^2/2$, where m is the particle mass and $\tilde{a}(\omega_0)$ is the Fourier transform of the acceleration profile $a(t)$ evaluated at the trap frequency ω_0. When $a(t)$ is linear in time, this energy depends

on the total displacement D and the time of movement T as approximately $\Delta E = 36mD^2/(\omega_0^2 T^4)$. Based on this estimate, it is reasonable to achieve substantial atom displacements $D > 50\,\mu$m within 250 μs for performing feed-forward applications: for typical trap frequencies $\omega_0 \approx 2\pi \times 50$ kHz, the atom's vibrational quantum number would increase by only $\Delta N < 1$. Indeed, such transport has been demonstrated by Ref. [16] without significant decoherence or atom loss due to heating. Moving the atoms by a distance D would then suppress reabsorption rates during ancilla readout to $\sigma/(4\pi D^2)$, where σ is the absorption cross-section [151]. Moreover, detuning the optical transitions for ancilla atoms by Δ further suppresses reabsorption by a factor $\approx (\Gamma/2\Delta)^2$, where Γ is the resonance linewidth, and $\Delta > 10\,\Gamma$ can be readily achieved with moderate powers of a light-shifting beam [151]. Between moving and light-shifting the ancillary atoms, cross-talk errors on the data qubits can be suppressed by five or more orders of magnitude, to negligible levels.

Alternatively, the measurement of ancilla qubit states can be achieved by using two different atomic species for the data and ancilla atoms (such as two different isotopes of the same atom or two different atomic species). In this approach [179,148], the ancilla atoms can still interact with the data atoms when both are coupled to Rydberg states, while they can be measured independently without disturbing the data atom states.

Finally, fast detection schemes were recently demonstrated in experiments with small atomic ensembles using Rydberg electromagnetically induced transparency (EIT) [174]. These could be potentially utilized to identify Rydberg population after entangling gates. These schemes could be incorporated into the tweezer array platforms by creating larger, elongated traps at selected locations containing optically dense atomic ensembles. In this approach, the Rydberg blockade effect leads to a sharp signature in the absorption spectrum of a weak EIT probe beam depending on whether a nearby Rydberg atom is present. Due to the collectively enhanced Rabi frequency, the detection time can be reduced to $\sim 6\mu$s [174], comparable to the duration of an entangling gate. This ultrafast, non-destructive Rydberg atom detector thus provides a promising implementation for the measure-

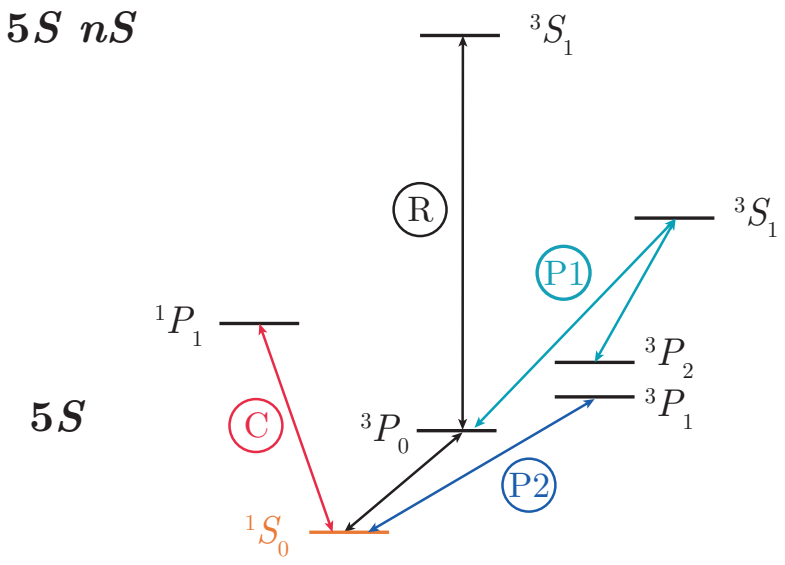


Figure 6.9: Relevant level diagram for implementing our FTQC protocols with neutral alkaline earth Rydberg atoms such as ^{87}Sr. The qubit is encoded in the stretched $^1S_0$ ground state (orange). Transitions to a $5SnS$, $^3S_1$ Rydberg state can be driven by first coherently mapping one of the qubit states to the $^3P_0$ clock state and then exciting the clock state to the Rydberg state (R, black). Optical pumping to correct for non-Rydberg leakage is implemented in two stages by driving the light blue transitions (P1) followed by the dark blue transition (P2). State readout and strong cooling for state initialization are implemented via the $^1S_0 \leftrightarrow {}^1P_1$ transition (C, red), while narrow-line cooling can be implemented via the P2 transition.

ment and feed-forward corrections needed for our protocols.

In this work, we have focused primarily on developing FTQC protocols for neutral alkali atoms coupled to Rydberg states. Recently, significant progress has also been made towards using alkaline earth(-like) atoms such as Sr and Yb for Rydberg-based quantum computations [110,170]. In this section, we show how our methods can also be applied for such setups. While we focus on an example of ^{87}Sr for concreteness, our discussion is generic for fermionic species of alkaline earth(-like) atoms.

For alkaline earth(-like) atoms, the $^1S_0$ ground states have no electronic orbital or spin angular momentum, so the only source of degeneracy is the nonzero nuclear spin (which can be quite large, e.g. $I = 9/2$ for ^{87}Sr). For our protocols, a most convenient qubit encoding uses the stretched ground states: $|0\rangle \equiv |m_I = -I\rangle, |1\rangle \equiv |m_I = +I\rangle$. In this encoding, strong cooling and

state readout can be implemented via the ${}^1S_0 \leftrightarrow {}^1P_1$ transition, while narrow-line cooling can be performed on the ${}^1S_0 \leftrightarrow {}^3P_1$ transition. Entangling gates can be implemented by selectively exciting the $|1\rangle$ state to a stretched Rydberg ${}^3S_1$ state. This state selectivity can be achieved by coherently mapping one of the qubit states to the ${}^3P_0$ clock state, performing Rydberg pulses between the clock state and the Rydberg state, and mapping back to the ${}^1S_0$ ground state, where we have utilized the linear Zeeman shift in the clock transition arising from hyperfine coupling between the ${}^3P_0$ and ${}^3P_1$ states[20]. The relevant level diagram is shown in Figure 6.9 for the case of ^{87}Sr (see also Ref. [108]).

During these entangling operations, an atom in the Rydberg state may undergo various errors such as BBR transitions, RD, or intermediate state scattering as described in Section 6.3. For alkaline earth(-like) atoms, the resulting Kraus operators can be described by Pauli-Z errors and quantum jumps to Rydberg states, ${}^1S_0$ ground states, or metastable ${}^3P$ states as allowed by dipole selection.

Following our approach for alkali atoms, we must convert all such errors to Pauli-Z errors to apply our FTQC protocols. By using ancilla atoms and the blockade effect, the quantum jumps to Rydberg states can be corrected in the same fashion as for alkali atoms. However, due to the presence of metastable ${}^3P$ levels, the correction of non-Rydberg leakage errors is more complicated, and the optical pumping must be done in two stages (see Figure 6.9): (1) Use σ^+-polarized light from the ${}^3P_{0,2}$ states to the triplet excited ${}^3S_1$ state to re-pump all ${}^3P$ states to the ${}^3P_1$ manifold; these states will decay back into the ${}^1S_0$ ground states. (2) Use σ^+-polarized on the narrow-line cooling transition ${}^1S_0 \leftrightarrow {}^3P_1$ to pump ground states with $m_F > -I$ to the stretched ground state $|1\rangle = |m_I = +I\rangle$. After these two steps, all non-Rydberg leakage errors will be mapped to the error $|1\rangle\langle 1|$, which is expressible in terms of Pauli-Z errors. We note that while Pauli-X errors could in principle arise from polarization impurities in the ${}^1S_0 \leftrightarrow {}^3P_1$ beam in the second stage, this would require several consecutive polarization imperfections, each of which has a very low probability of roughly $0.2 - 0.5\%$; thus, the overall probability of Pauli-X errors arising from imperfect polar-

ization is negligible. Therefore, by using this optical pumping scheme to convert all non-Rydberg leakage errors to Z errors, the FTQC schemes of Sections 6.4 and 6.5 can be implemented in alkaline earth(-like) atoms.

6.7 Conclusions

In this work, we have presented a detailed analysis of the dominant error channels arising in quantum computation using neutral Rydberg atoms. We show that although the multilevel nature of atoms and the complex decay channels for Rydberg states lead to many additional types of errors not considered in traditional QEC settings, the specific structure of the error model allows us to design hardware-efficient FTQC protocols based on the seven-qubit and hardware-tailored three-qubit codes with significantly reduced overhead compared to general-purpose schemes. The crux of these results is the ability to convert the complicated error model to Pauli-Z errors by introducing ancilla atoms and making use of the Rydberg blockade effect, dipole selection rules, and new schemes for optical pumping. To use the three-atom repetition code, we designed a new laser pulse sequence to implement bias-preserving CNOT and Toffoli gates. For both protocols, we propose simple, scalable geometrical layouts and demonstrate feasibility of all components of FTQC for near-term implementation.

Our results provide an important step towards building large-scale quantum computers based on neutral atom setups and highlight the importance of designing FTQC schemes based on the specific structure of the error model and the unique capabilities of the hardware platform. Compared to some general-purpose FTQC protocols, our hardware-efficient approaches for Rydberg systems enable an order-of-magnitude improvement in resource overhead in terms of the number of physical gates or required ancillas. We believe many of the ideas developed in this work, such as the exploitation of the multi-level structure of the physical system, are transferable to other quantum

computing platforms such as trapped ions and superconducting qubits. In the former case, an optical pumping-based protocol to mitigate leakage in ions with low nuclear spin such as $^{171}Yb^{+}$ was recently developed and realized experimentally[76]; we believe that insights from our work would be helpful for developing more general leakage correction methods in such setups and incorporating them into FTQC protocols.

It matters not how strait the gate,

How charged with punishments the scroll,

I am the master of my fate:

I am the captain of my soul.

William Ernest Henley

7

Conclusions and Future Directions

In conclusion, we have presented several solutions to two important questions for quantum information theory, as introduced in Section 1.2:

1. Given the apparent difficulty of classically simulating large quantum systems, are there any useful applications in which current or near-term quantum devices can outperform classical computers?

2. Can we develop efficient error-correction schemes to suppress the error rates and further improve qubit and gate fidelities in near-term quantum devices?

We focused on the problem of quantum phase recognition in answering the first question; meanwhile, we approached the second using the concept of hardware-efficient quantum error correction. More specifically, we developed a novel quantum machine learning paradigm in Chapter 4, the quantum convolutional neural network (QCNN), and we demonstrated its success at both recognizing one-dimensional quantum phases and optimizing quantum error-correcting codes. Drawing inspiration from this QCNN framework, we then introduced in Chapter 5 a novel method (LED) to detect and characterize two-dimensional topological phases of matter by using techniques from quantum error correction. Through theoretical proofs, numerical simulations, and an experimental realization, we found that LED is substantially more efficient and robust than traditional approaches to recognize topological order. Finally, in Chapter 6, we developed hardware-efficient, fault-tolerant protocols to perform quantum computation with neutral Rydberg atoms. By leveraging the unique advantages of Rydberg atoms to protect against the hardware-specific error model, our protocols became substantially more efficient than existing, general-purpose approaches.

Following these exciting results, there are multiple intriguing future directions to consider for each of these works, which we outline below.

7.1 Future Directions for QCNN

In Chapter 4, we found that QCNNs provide a promising quantum machine learning paradigm. Several interesting generalizations and future applications can be considered. For example, to recognize more exotic phases, we could also relax the translation-invariance constraints, resulting in $O(N \log N)$ parameters for system size N, or we can use ancilla qubits to implement parallel feature maps following traditional CNN architecture. Further extensions can incorporate optimizations

for fault-tolerant operations on quantum error-correcting code spaces: for instance, it would be particularly exciting if one could use quantum machine learning ideas to develop hardware-efficient, fault-tolerant quantum computing protocols such as the ones we designed in Chapter 6. In addition, while we have used a finite-difference scheme to compute gradients in our learning demonstrations, the structural similarity of QCNN with its classical counterpart motivates adoption of more efficient schemes such as backpropagation[99].

Recently, our work on QCNN has seen two exciting developments. First, the application of QCNNs for quantum phase recognition was recently demonstrated on a seven-qubit superconducting quantum processor[77]. Second, the QCNN has recently been featured as a major example on Google's new TensorFlow Quantum package[25], enabling a wide range of quantum enthusiasts to further explore our framework. It will be particularly interesting to see how these developments can lead to larger-scale experimental realizations or the discovery of novel applications for QCNNs.

7.2 Future Directions for LED

Our results in Chapter 5 demonstrate that our LED method constitutes a promising approach to enhance the detection and characterization of topological order. Several generalizations and future avenues can be considered. For example, while our present LED work analyzes a spin-liquid state prepared using a Rydberg-atom quantum simulator, LED is also directly applicable to other quantum simulation and computation platforms, such as superconducting qubits[138] or trapped ions[153]. Moreover, while our current analysis uses classical post-processing of Z- and X-basis experimental snapshots, it can be mapped to a quantum circuit model within the QCNN framework of Chapter 4. Such a quantum circuit can allow for simultaneous measurement of both Z- and X-basis decorated loop operators in each experimental repetition. Furthermore, the variational framework of QCNN circuits can enable adaptive measurement procedures, which can better accommodate states

differing from the fixed-point by large quasi-local unitary rotations. Another promising direction is to study the relationship between the "correctability" of states in our mixed-state phase diagram (Figure 5.4c) and other methods of studying topological order in mixed states, such as entanglement negativity [125,79,101]. In particular, it could be intriguing to further explore the dependence of the correctable regime on the choice of local error correction and/or coarse-graining procedure. Finally, while our approach can be directly applied to any abelian topological phase described by the quantum double of a finite abelian group [92], it could be interesting to consider more exotic topological phases, such as fracton phases, non-abelian topological orders, or gauge theories with continuous gauge groups [162,163]. Such methods can then become indispensable parts of quantum simulation toolboxes for understanding exotic states of entangled quantum matter.

7.3 Future Directions for Fault-Tolerant Quantum Computation with Rydberg Atoms

The hardware-efficient, fault-tolerant quantum computation protocols which we developed in Chapter 6 provide an important step towards building large-scale quantum computers based on neutral atom setups. Several interesting extensions can be considered. For example, while we have primarily quantified the resource cost for fault-tolerant quantum computation proposals by studying the number of qubits and gates required, another related and commonly used metric is the error threshold, which amounts to the physical qubit and gate fidelities required to produce logical error rates that are lower than the physical error rate. One may estimate the error thresholds for two- and three-qubit gates directly by using the numbers presented in Tables 6.1 and 6.2 and requiring each logical operation or stabilizer measurement step to have at most a single error, but it can be more useful to obtain the precise numbers for these thresholds via numerical simulation.

A more detailed study of the error threshold would be especially helpful if one intends to extend

our work to codes with distance greater than 3 and compare the relative performance and scalability of these approaches with our current proposal. One particularly intriguing direction could be to evaluate the performance and resource cost of existing fault-tolerant quantum computation protocols based on topological codes such as surface codes or color codes [22,58,18,19,155,4] upon applying our techniques to address Rydberg and non-Rydberg leakage errors. Indeed, one recent work has already demonstrated ultrahigh error threshold in the surface code when the underlying noise model is biased [156]; this motivates the use of surface codes in a Rydberg system where dominant errors have all been converted to Pauli-Z type. Going further in this direction, one could eliminate all of the Rydberg-specific leakage errors using our fault-tolerant quantum computation protocols, and then concatenate our codes with more traditional quantum error correction approaches to address any higher-order Pauli-X or Y-type errors that were neglected in our studies, or to further suppress the logical error rate to even higher orders.

Experimentally, there have been many rapid developments in Rydberg-atom quantum computing in the last few years, and further advances are expected in the near future. In particular, the newly-demonstrated ability to transport atomic qubits while preserving quantum coherence opens many doors for performing error correction and fault-tolerant quantum computing [16]. It would be particularly exciting to combine this groundbreaking new feature of neutral-atom platforms with our hardware-efficient protocols to enable large-scale fault-tolerant quantum computation with Rydberg atoms.

A

Supplementary Information for Chapter 4

A.1 Phase Diagram and QCNN Circuit Simulations

The phase diagram in Chapter 4 (Figure 4.2a) was numerically obtained using the infinite size density-matrix renormalization group (DMRG) algorithm. We generally follow the method outlined in Ref. [114] with the maximum bond dimension 150. To extract each data point in Figure 4.2a, we numerically obtain the ground state energy density as a function of h_2 for fixed h_1 and computed

its second order derivative. The phase boundary points are identified from sharp peaks.

The simulation of our QCNN in Figure 4.2b also utilizes matrix product state representations. We first obtain the input ground state wavefunction using finite-size DMRG [114] with bond dimension $D = 130$ for a system of $N = 135$ qubits. Then, the circuit operations are performed by sequentially applying SWAP and two-qubit gates on nearest neighboring qubits [164]. Each three-qubit gate is decomposed into two-qubit unitaries [122]. We find that increasing bond dimension to $D = 150$ does not lead to any visible changes in our main figures, confirming a reasonable convergence of our method. The color plot in Figure 4.2a is similarly generated for a system of $N = 45$ qubits.

A.2 Multiscale String Order Parameters

We examine the final operator measured by our circuit that recognizes the SPT phase in the Heisenberg picture. Although a QCNN performs non-unitary measurements in the pooling layers, similar to QEC circuits [129], one can postpone all measurements to the end and replace pooling layers by unitary controlled gates acting on both measured and unmeasured qubits. In this way, the QCNN is equivalent to measuring a non-local observable

$$\mathcal{O} = (U_{\mathrm{CP}}^{(d)} \ldots U_{\mathrm{CP}}^{(1)})^{\dagger} Z_{i-1} X_i Z_{i+1} (U_{\mathrm{CP}}^{(d)} \ldots U_{\mathrm{CP}}^{(1)}) \tag{A.1}$$

where i is the index of the measured qubit in the final layer and $U_{\mathrm{CP}}^{(l)}$ is the unitary corresponding to the convolution-pooling unit at depth l. A more explicit expression of $\mathcal{O}$ can be obtained by commuting U_{CP} with the Pauli operators, which yields recursive relations:

$$U_{\mathrm{CP}}^{\dagger} X_i U_{\mathrm{CP}} = X_{\tilde{i}-2} X_{\tilde{i}} X_{\tilde{i}+2} \tag{A.2}$$

$$U_{\mathrm{CP}}^{\dagger} Z_i U_{\mathrm{CP}} = \frac{1}{2}(Z_{\tilde{i}} + Z_{\tilde{i}-2}X_{\tilde{i}-1} + X_{\tilde{i}+1}Z_{\tilde{i}+2} - Z_{\tilde{i}-2}X_{\tilde{i}-1}Z_{\tilde{i}}X_{\tilde{i}+1}Z_{\tilde{i}+2}) \tag{A.3}$$

$\tilde{i}$ enumerates every qubit at depth $l-1$, including those measured in the pooling layer. It follows that an SOP of the form $ZXX...XZ$ at depth l transforms into a weighted linear combination of 16 products of SOPs at depth $l-1$. Thus, instead of measuring a single SOP, our QCNN circuit measures a sum of products of exponentially many different SOPs:

$$\mathcal{O} = \sum_{ab} C_{ab}^{(1)} \mathcal{S}_{ab} + \sum_{a_1 b_1 a_2 b_2} C_{a_1 b_1 a_2 b_2}^{(2)} \mathcal{S}_{a_1 b_1} \mathcal{S}_{a_2 b_2} + \cdots , \tag{A.4}$$

$\mathcal{O}$ can be viewed as a *multiscale* string order parameter with coefficients computed recursively in d using Eqs. (A.2,A.3). This allows the QCNN to produce a sharp classification output even when the correlation length is as long as 3^d.

A.3 Construction of QCNN Circuit

To construct the exact QCNN circuit in Figure 4.2b, we followed the guidelines discussed in Section 4.3. Specifically, we designed the convolution and pooling layers to satisfy the following two important properties:

1. Fixed-point criterion: If the input is a cluster state $|\psi_0\rangle$ of L spins, the output of the convolution-pooling layers is a cluster state $|\psi_0\rangle$ of $L/3$ spins, with all measurements deterministically yielding $|0\rangle$.

2. QEC criterion: If the input is not $|\psi_0\rangle$ but instead differs from $|\psi_0\rangle$ at one site by an error which commutes with the global symmetry, the output should still be a cluster state of $L/3$

spins, but at least one of the measurements will result in the state $|1\rangle$.

These two properties are desirable for any quantum circuit implementation of RG flow for performing QPR.

In the specific case of our Hamiltonian, the ground state (1D cluster state) is a graph state, which can be efficiently obtained by applying a sequence of controlled phase gates to a product state. This significantly simplifies the construction of the MERA representation for the fixed-point criterion. To satisfy the QEC criterion, we treat the ground state manifold of the unperturbed Hamiltonian $H = -J\sum_i Z_i X_{i+1} Z_{i+2}$ as the code space of a stabilizer code with stabilizers $\{Z_i X_{i+1} Z_{i+2}\}$. The remaining degrees of freedom in the QCNN convolution and pooling layers are then specified such that the circuit detects and corrects the error (i.e. measures at least one $|1\rangle$ and prevents propagation to the next layer) when a single-qubit X error is present.

A.4 QCNN for General QPR Problems

Our interpretation of QCNNs in terms of MERA and QEC motivates their application for recognizing more generic quantum phases. For any quantum phase $\mathcal{P}$ whose RG fixed-point wavefunction $|\psi_0(\mathcal{P})\rangle$ has a tensor network representation in isometric or G-isometric form [139] (Figure A.1a), one can systematically construct a corresponding QCNN circuit. This family of quantum phases includes all 1D SPT and 2D string-net phases [41,139,103]. In these cases, one can explicitly construct a commuting parent Hamiltonian for $|\psi_0(\mathcal{P})\rangle$ and a MERA structure in which $|\psi_0(\mathcal{P})\rangle$ is a fixed-point wavefunction (Figure A.1a for 1D systems). The diagrammatic proof of this fixed-point property is given in Figure A.1b. Furthermore, any "local error" perturbing an input state away from $|\psi_0(\mathcal{P})\rangle$ can be identified by measuring a fraction of terms in the parent Hamiltonian, similar to syndrome measurements in stabilizer-based QEC [129]. Then, a QCNN for $\mathcal{P}$ simply consists of the MERA for $|\psi_0(\mathcal{P})\rangle$ and a nested QEC scheme in which an input state with error density below the

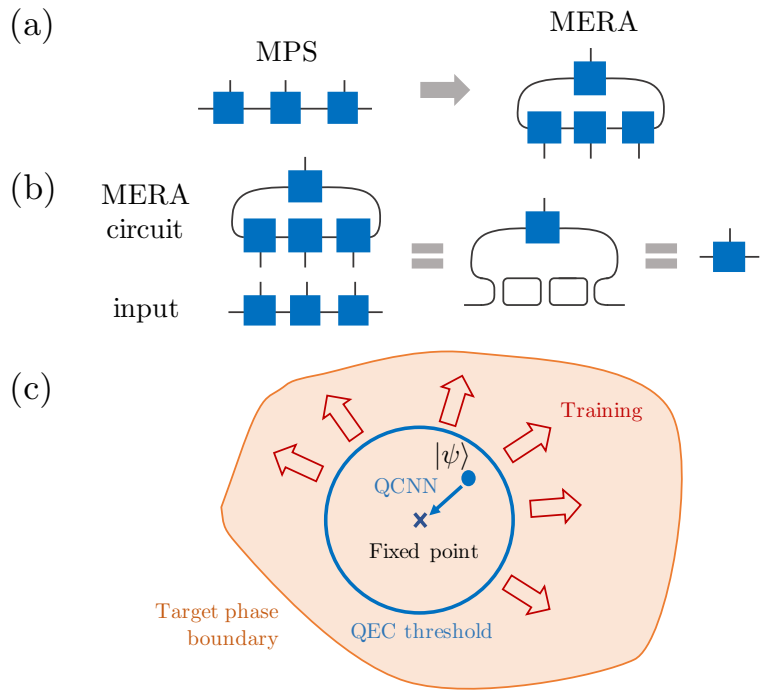


Figure A.1: QCNN for generic QPR problems. (a) Given a state with a translationally invariant, isometric matrix product state representation (e.g. a fixed point state for a 1D SPT phase), we explicitly construct an isometry for the MERA representation of this state. Blue squares are the matrix product state tensors, while black lines are the legs of the tensor. While we have illustrated a 3-to-1 isometry, the generalization to arbitrary n-to-1 isometries is straightforward. (b) Diagrammatic proof showing that a MERA constructed from the above tensor maps the fixed-point state back to a shorter version of itself. The first equality uses the definition of isometric tensor, and loops in the middle diagram simplify to a constant number unity. The generalization of this isometry to higher dimensions is discussed in Ref. 56. (c) One helpful initial parameterization for QPR problems consists of a MERA for the fixed point state $|\psi_0(\mathcal{P})\rangle$ and a choice of nested QEC, so that states within the QEC threshold flow toward $|\psi_0(\mathcal{P})\rangle$. Training procedures then expand this threshold boundary to the phase boundary.

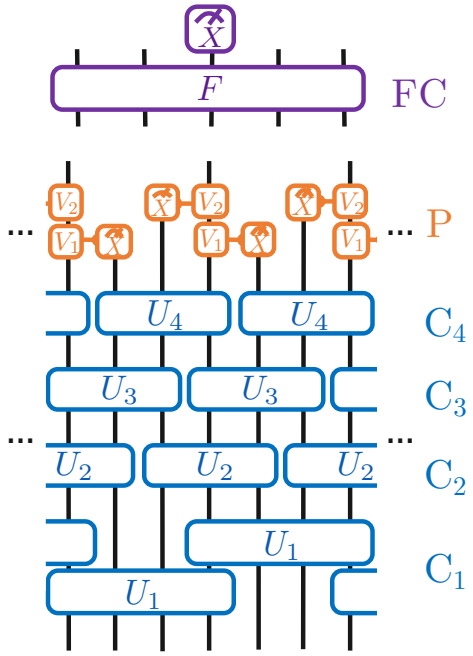


Figure A.2: Circuit parameterization for training a QCNN to solve QPR. Our circuit involves 4 different convolution layers ($C_1 - C_4$), a pooling layer, and a fully connected layer. The unitaries are initialized to random values, and learned via gradient descent.

QEC threshold[2] "flows" to the RG fixed point. Such a QCNN can be optimized via our learning procedure.

While our generic learning protocol begins with completely random unitaries, as in the classical case[99], this initialization may not be the most efficient for gradient descent. Instead, motivated by deep learning techniques such as pre-training[99], a better initial parameterization would consist of a MERA representation of $|\psi_0(\mathcal{P})\rangle$ and one choice of nested QEC. With such an initialization, the learning procedure serves to optimize the QEC scheme, expanding its threshold to the target phase boundary (Figure A.1c).

A.5 Demonstration of Learning Procedure for QPR

To perform our learning procedure in a QPR problem, we choose the hyperparameters for the QCNN as shown in Figure A.2. This hyperparameter structure can be used for generic (1D) phases, and is characterized by a single integer n that determines the reduction of system size in each convolution-pooling layer, $L \to L/n$. (Figure A.2 shows the special case where $n = 3$). The first convolution layer involves $(n+1)$-qubit unitaries starting on every n^{th} qubit. This is followed by n layers of n-qubit unitaries arranged as shown in Figure A.2. The pooling layer measures $n-1$ out of every contiguous block of n qubits; each of these is associated with a unitary V_j applied to the remaining qubit, depending on the measurement outcome. This set of convolution and pooling layers is repeated d times, where d is the QCNN depth. Finally, the fully connected layer consists of an arbitrary unitary on the remaining N/n^d qubits, and the classification output is given by the measurement output of the middle qubit (or any fixed choice of one of them). For our example, we choose $n = 3$ because the Hamiltonian in Equation 4.2 involves three-qubit terms.

In our simulations, we consider only $N = 15$ spins and depth $d = 1$, because simulating quantum circuits on classical computers requires a large amount of resources. We parameterize unitaries as exponentials of generalized $a \times a$ Gell-Mann matrices $\{\Lambda_i\}$, where $a = 2^w$ and w is the number of qubits involved in the unitary [13]: $U = \exp\left(-i\sum_j c_j \Lambda_j\right)$.

This parameterization is used directly for the unitaries in the convolution layers $C_2 - C_4$, the pooling layer, and the fully connected layer. For the first convolution layer C_1, we restrict the choice of U_1 to a product of six two-qubit unitaries between each possible pair of qubits:

$$U_1 = U_{(23)}U_{(24)}U_{(13)}U_{(14)}U_{(12)}U_{(34)}, \tag{A.5}$$

where $U_{(\alpha\beta)}$ is a two-qubit unitary acting on qubits indexed by α and β. Such a decomposition is

useful when considering experimental implementation.

In the QCNN learning procedure, all parameters c_μ are set to random values between 0 and 2π for the unitaries $\{U_i, V_j, F\}$. In every iteration of gradient descent, we compute the derivative of the mean-squared error function (Equation 4.1 in Section 4.2) to first order with respect to each of these coefficients c_μ by using the finite-difference method:

$$\frac{\partial \text{MSE}}{\partial c_\mu} = \frac{1}{2\varepsilon}\left(\text{MSE}(c_\mu + \varepsilon) - \text{MSE}(c_\mu - \varepsilon)\right) + O(\varepsilon^2). \tag{A.6}$$

Each coefficient is thus updated as $c_\mu \mapsto c_\mu - \eta \frac{\partial \text{MSE}}{\partial c_\mu}$, where η is the learning rate for that iteration. We compute the learning rate using the bold driver technique from machine learning, where η is increased by 5% if the error has decreased from the previous iteration, and decreased by 50% otherwise[78]. We repeat the gradient descent procedure until the error function changes on the order of 10^{-5} between successive iterations. In our simulations, we use $\varepsilon = 10^{-4}$ for the gradient computation, and begin with an initial learning rate of $\eta_0 = 10$.

A.6 Demonstration of Learning Procedure for QEC

To obtain the QEC code considered in Section 4.4, we consider a QCNN with $N = 9$ input physical qubits and simulate the circuit evolution of its $2^N \times 2^N$ density matrix exactly. Strictly speaking, our QCNN has three layers: a three-qubit convolution layer U_1, a 3-to-1 pooling layer, and a 3-to-1 fully connected layer U_2. Without loss of generality, we may ignore the optimization over the pooling layer by absorbing its effect into the first convolution layer, leading to the effective two-layer structure shown in Figure 4.6a. The generic three-qubit unitary operations U_1 and U_2 are parameterized using 63 Gell-Mann coefficients each.

As discussed in Section 4.4, we consider three different error models: (1) independent single-qubit errors on all qubits with equal probabilities p_μ for $\mu = X$, Y, and Z errors, (2) independent

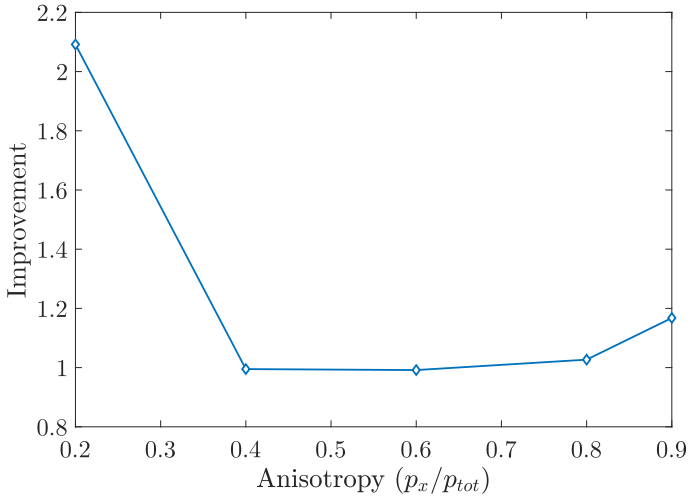


Figure A.3: Ratio between the logical error rate of the Shor code and that of the QCNN code for the anisotropic depolarization error model. We fix the total input error rate $p_{\text{tot}} = p_x + p_y + p_z = 0.001$ and $p_y = p_z$, while varying the ratio p_x/p_{tot}.

single-qubit errors on all qubits, with anisotropic probabilities $p_x \neq p_y = p_z$, and (3) independent single-qubit anisotropic errors with additional two-qubit correlated errors $X_i X_{i+1}$ with probability p_{xx}. More specifically, the first two error models are realized by applying a (generally anisotropic) depolarization quantum channel to each of the nine physical qubits:

$$\mathcal{N}_{1,i} : \rho \mapsto (1 - \sum_{\mu} p_\mu)\rho + \sum_{\mu} p_\mu \sigma_i^\mu \rho \sigma_i^\mu \tag{A.7}$$

with Pauli matrices σ_i^μ for $i \in \{1, 2, \ldots, 9\}$ (the qubit indices are defined from bottom to top in Figure 4.6a). For the anisotropic case, we trained the QCNN on various different error models with the same total error probability $p_x + p_y + p_z = 0.001$, but different relative ratios; the resulting ratio between the logical error probability of the Shor code and that of the QCNN code is plotted as a function of anisotropy in Figure A.3. For strongly anisotropic models, the QCNN outperforms the Shor code, while for nearly isotropic models, the Shor code is optimal and QCNN can achieve

the same logical error rate.

For the correlated error model, we additionally apply a quantum channel:

$$\mathcal{N}_{2,i} : \rho \mapsto (1 - p_{xx})\rho + p_{xx} X_i X_{i+1} \rho X_i X_{i+1} \tag{A.8}$$

for pairs of nearby qubits, i.e. $i \in \{1, 2, 4, 5, 7, 8\}$. Such a geometrically local correlation is motivated from experimental considerations. In this case, we train our QCNN circuit on a specific error model with parameter choices $p_x = 5.8 \times 10^{-3}, p_y = p_z = 2 \times 10^{-3}, p_{xx} = 2 \times 10^{-4}$ and evaluate the logical error probabilities for various physical error models with the same relative ratios, but different total error per qubit $p_x + p_y + p_z + p_{xx}$. In general, for an anisotropic logical error model with probabilities p_μ for σ_μ logical errors, the overlap f_q is $(1 - 2\sum_\mu p_\mu/3)$, since $\langle \pm\nu | \sigma_\mu | \pm\nu \rangle = (-1)^{\delta_{\mu,\nu}+1}$. Becuase of this, we compute the total logical error probability from f_q as $1.5(1 - f_q)$. Hence, our goal is to maximize the logical state overlap f_q defined in Equation (4.6). If we naively apply the gradient descent method based on f_q directly to both U_1 and U_2, we find that the optimization is easily trapped in a local optimum. Instead, we optimize two unitaries U_1 and U_2 sequentially, similar to the layer-by-layer optimization in backpropagation for conventional CNN[99].

A few remarks are in order. First, since U_1 is optimized prior to U_2, one needs to devise an efficient cost function C_1 that is independent of U_2. In particular, simply maximizing f_q with an assumption $U_2 = 1$ may not be ideal, since such choice does not capture a potential interplay between U_1 and U_2. Second, because U_1 captures arbitrary single qubit rotations, the definition of C_1 should be basis independent. Finally, we note that the tree structure of our circuit allows one to view the first layer as an independent quantum channel:

$$\mathcal{M}_{U_1} : \rho \mapsto \mathrm{tr}_a[U_1 \mathcal{N}(U_1^\dagger(|0\rangle\langle 0| \otimes \rho \otimes |0\rangle\langle 0|)U_1)U_1^\dagger], \tag{A.9}$$

where $\mathrm{tr}_a[\cdot]$ denotes tracing over the ancilla qubits that are measured in the intermediate step. From this perspective, $\mathcal{M}_{U_1}$ describes an effective error model to be corrected by the second layer.

With these considerations, we optimize U_1 such that the effective error model $\mathcal{M}_{U_1}$ becomes as classical as possible, i.e. $\mathcal{M}_{U_1}$ is dominated by a "flip" error along a certain axis with a strongly suppressed "phase" error. Only then, the remant, simpler errors will be corrected by the second layer. More specifically, one may represent $\mathcal{M}_{U_1}$ using a map $M_{U_1} : \mathbf{r} \mapsto M\mathbf{r} + \mathbf{c}$, where $\mathbf{r} \in \mathbb{R}^3$ is the Bloch vector for a qubit state $\rho \equiv \frac{1}{2}1 + \mathbf{r} \cdot$ [122]. The singular values of the real matrix M encode the probabilities $p_1 \geq p_2 \geq p_3$ for three different types of errors. We choose our cost function for the first layer as $C_1 = p_1^2 + p_2 + p_3$, which is relatively more sensitive to p_2 and p_3 than p_1 and ensure that the resultant, optimized channel $\mathcal{M}_{U_1}$ is dominated by one type of error (with probability p_1). We note that M can be efficiently evaluated from a quantum device without knowing $\mathcal{N}$, by performing quantum process tomography for a single logical qubit. Once U_1 is optimized, we use gradient decent to find an optimal U_2 to maximize the fidelity f_q. As with QPR, gradients are computed via the finite-difference method, and the learning rate is determined by the bold driver technique[99].

A.7 Experimental Resource Analysis

To compute the gate depth of the cluster model QCNN circuit in a Rydberg atom implementation, we analyze each gate shown in Figure 4.2b. By postponing pooling layer measurements to the end of the circuit, the multi-qubit gates required are

$$\mathrm{C}_z\mathrm{Z}_{ij} = e^{i\pi(-1+Z_i)(-1+Z_j)/4} \tag{A.10}$$

$$\mathrm{C}_x\mathrm{Z}_{ij} = e^{i\pi(-1+X_i)(-1+Z_j)/4} \tag{A.11}$$

$$C_x C_x X_{ijk} = e^{i\pi(-1+X_i)(-1+X_j)(-1+X_k)/8}. \tag{A.12}$$

By using Rydberg blockade-mediated controlled gates [136], it is straightforward to implement $C_z Z_{ij}$ and $C_z C_z Z_{ijk} = e^{i\pi(-1+Z_i)(-1+Z_j)(-1+Z_k)/8}$. The desired $C_x Z_{ij}$ and $C_x C_x X_{ijk}$ gates can then be obtained by conjugating $C_z Z_{ij}$ and $C_z C_z Z_{ijk}$ by single-qubit rotations. For input size of N spins, the k^{th} convolution-pooling unit thus applies $4N/3^{k-1}$ $C_z Z_{ij}$ gates, $N/3^{k-1}$ $C_x C_x X_{ijk}$ gates, and $2N/3^{k-1}$ layers of $C_x Z_{ij}$ gates. The depth of single-qubit rotations required is $4d$, as these rotations can be implemented in parallel on all N qubits. Finally, the fully connected layer consists of $N3^{1-d}$ $C_z Z_{ij}$ gates. Thus, the total number of multi-qubit operations required for a QCNN of depth d operating on N spins is $\frac{7N}{2}(1 - 3^{1-d}) + N3^{1-d}$. Note that we need not use SWAP gates since the Rydberg interaction is long-range.

B

Supplementary Information for Chapter 5

B.1 Numerical Simulations for the Toric Code

In this section, we explain how the numerical simulations underlying Figures 5.2, 5.3, and 5.4 are performed. We begin by constructing a projected entangled pair state (PEPS) representation of the

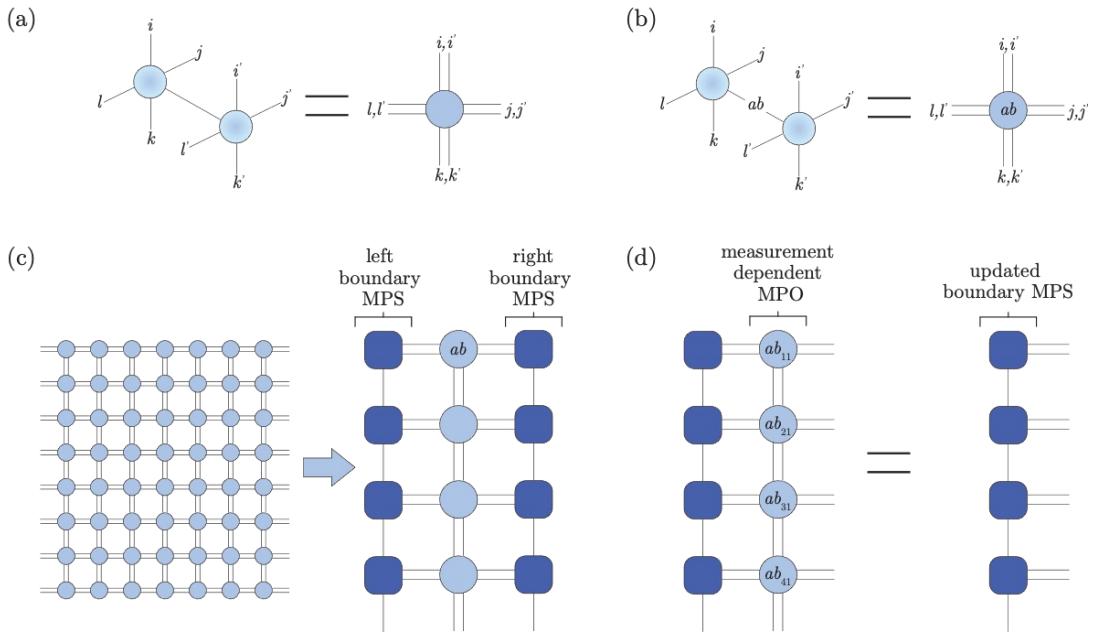


Figure B.1: PEPS sampling algorithm. Expectation values are computing with respect to both $|\psi\rangle$ (back) and $\langle\psi|$ (front). (a) Tracing, or averaging over measurement outcomes can be done by contracting the physical indices, and is needed to compute marginal probabilities. (b) To compute the probability of a particular Z basis measurement, the physical index is assigned a particular value ab. (c) We can efficiently contract a 2D PEPS tensor network on an infinite strip of finite height, by using a left and right boundary MPS (only top four rows shown). The probability distribution for projective measurements on a particular site, e.g. $xy = 11$ can then be computed efficiently. (d) Once an entire column has been sampled, the measurement-dependent MPO can applied to the boundary MPS. Although performing this contraction exactly causes the bond-dimension to grow rapidly, away from phase boundaries, finite bond dimension is sufficient for accurate simulation. See Appendix B.9 for more details.

exact toric code ground state [140]. This construction utilizes a parity tensor P defined as

$$P_{ijkl} = \begin{cases} 1 & \text{if } i+j+k+l = 0 \mod 2 \\ 0 & \text{otherwise} \end{cases} \tag{B.1}$$

where each index $i, j, k, l \in \{0, 1\}$ (i.e., the tensor P has bond dimension two). Because the toric code is defined with qubits on the links of a square lattice, our PEPS representation of the state has one PEPS tensor with two physical indices per unit cell. Letting p, q be the physical indices and $ijkl$ be the virtual indices, the toric code PEPS tensor A is then given by $A^{pq}_{ijkl} = \delta^p_i \delta^q_j P_{ijkl}$. Our perturbed states $|\psi(g_X, g_Z)\rangle$ are constructed from the toric code state by applying imaginary time evolution to each site $L(g_X, g_Z) = e^{g_X X + g_Z Z}$:

$$A(g_X, g_Z)^{pq}_{ijkl} = \sum_{p',q'} L(g_X, g_Z)^p_{p'} L(g_X, g_Z)^q_{q'} A^{p'q'}_{ijkl}. \tag{B.2}$$

Notice that this operation does not change the PEPS bond dimension, thereby allowing for efficient simulation.

Our goal is to simulate projective Z-basis measurements to serve as the "experimental snapshot" input in Figure 5.1b. The key ingredient which enables efficient sampling is an algorithm for efficiently computing marginal and conditional probabilities. To illustrate how the algorithm works, consider the following procedure. We first label every unit cell by its coordinate (x, y). There are four possible measurement outcomes at each unit cell, and we compute the probability $P(\sigma_{(1,1)} = ab)$ that measurement of the first site $(x, y) = (1, 1)$ yields the outcome $ab = 00, 01, 10,$ or 11. Next, we select a sample ab_{11} based on this probability distribution. With this information, we compute the probability distribution on the second site, $P(\sigma_{21} = ab | ab_{11})$ conditioned on the first measurement outcome, and sample the second measurement outcome ab_{21}. The process then repeats,

with each subsequent distribution being conditioned on all prior measurements.

Computing the probabilities requires contracting a 2D tensor network (Figure B.1), which is in general $\#P$-hard [141]. In practice, however, the states we encounter have finite correlation length, and the computation becomes remarkably efficient throughout much of the phase diagram [43]. In particular, we work on a strip of finite height L_x and infinite length L_y, and introduce boundary matrix product states (MPS) to efficiently capture the effect of the environment—that is, the sites different from the one currently being sampled [120]. Because singular-value decomposition (SVD) truncation is used at each step to prevent the bond dimension of the boundary MPS from growing exponentially [166], the method is approximate; however, we only discard singular values smaller than $< 10^{-8}$, so truncation errors are insignificant. Details of the boundary conditions and contraction ordering are discussed in Appendix B.8.

In our simulations, we choose $L_x = 300$ unit cells and sample 1000 columns, giving us access to very large snapshots with 600,000 qubits. To minimize boundary effects, we compute observables which are supported on sites at least 30 unit cells away from the boundaries. Near the phase boundaries, the bond dimension (entanglement) of the boundary MPS becomes large due to the large correlation length, which increases the computational demands for sampling (gray data points in Figure 5.3). We numerically confirm this phase boundary with an independent calculation (see Appendix B.9).

B.2 Details on Error-Correction and Coarse-Graining Procedures

In this section, we explain the details of the LED decoding and coarse-graining procedures and demonstrate how bare Wilson loops become decorated under the LED protocol. Without loss of generality, we consider Z-basis measurements, from which we can calculate plaquette stabilizers B_u. Here, each plaquette is labelled by the 2D coordinate of its unit cell $u = (x, y)$. Since there are two

qubits per unit cell, each qubit carries a coordinate and a link label v or h, depending on whether its corresponding edge in the square lattice is vertical or horizontal, respectively. Finally, the projective measurement outcomes are denoted by $\sigma \in \{+1, -1\}$ (see Figure B.2).

To illustrate local error correction, we consider the "pairing decoder," which flips a qubit if and only if its two neighboring plaquettes are simultaneously occupied. Importantly, to preserve locality, we first compute all stabilizer values, and then flip qubits based on these values. The decision of whether to flip any qubit then depends only on its value, and the values of the six adjacent qubits with which it shares a plaquette. Equivalently, this error correction procedure corresponds to an operator transformation

$$\sigma_{u+\hat{x},v} \to \sigma_{u+\hat{x},v} \left(1 + B_u + B_{u+\hat{x}} - B_u B_{u+\hat{x}}\right)/2 \tag{B.3}$$

$$\sigma_{u+\hat{y},h} \to \sigma_{u+\hat{y},v} \left(1 + B_u + B_{u+\hat{y}} - B_u B_{u+\hat{y}}\right)/2 \tag{B.4}$$

To ensure all local errors are removed after a finite number of LED steps, we also pair anyons which occupy two plaquettes separated by a diagonal, such as B_u and $B_{u+\hat{x}+\hat{y}}$. The locality of the decoder ensures that the support of any local operator only grows by a finite amount with each step. Subsequently, the coarse-graining procedure replaces each $b \times b$ block of plaquettes with a single plaquette whose value is the product of b^2 plaquettes; microscopically, this can be done by defining new qubits as a product of b corresponding qubits in the original lattice. The combination of a local pairing step and a coarse-graining step forms a layer of real-space RG; upon applying this RG procedure n times, one can correct errors of higher and higher weight.

The bare Wilson loops measured in the final state are equivalent to decorated loop operators acting on the original state. These decorated operators can be efficiently computed from projective measurement data, since the eigenstates of the LED Z-loop and X-loop operators are product states in the Z and X bases, respectively. Furthermore, in the operator transformation picture, any loop or

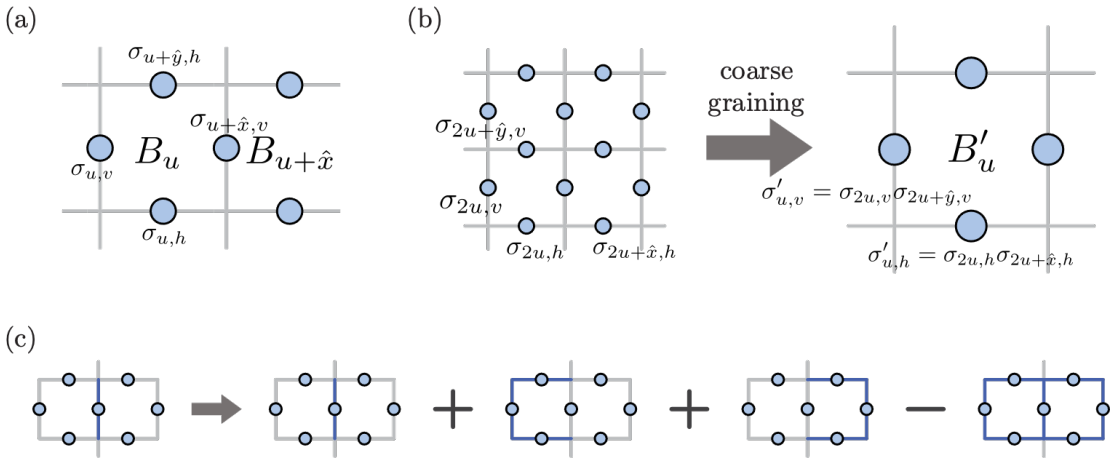


Figure B.2: LED coarse-graining and operator transformation. (a) In the toric code model, qubits are located on the links of a square lattice, and the stabilizer associated with any plaquette is given by a product of four single-qubit Pauli-Z operators. (b) Coarse-graining maps a $b \times b$ block of plaquettes to a single plaquette whose value is the product of the b^2 plaquettes (here $b = 2$). Microscopically, coarse-grained qubits σ' are products of b lower-level qubits σ. Coarse-grained stabilizers B'_u are therefore equivalent to a product of $b \times b$ stabilizers at the lower level. (c) Pairing correction flips a qubit conditioned on the state of its two neighboring stabilizers. This is equivalent to an operator transformation where the qubit is decorated by products of closed loops.

string of length L maps onto a linear combination of exponentially many ($2^{O(L)}$) loops or strings, respectively. Thus, while the operator transformation picture is helpful for conceptual reasons, it is computationally much easier to use the original picture of error-correction and coarse-graining.

A few remarks are in order. First, one important property of LED is that it preserves commutation relations: consider two anti-commuting X and Z strings which intersect at a single point, far from the strings' endpoints. Upon applying LED, the resulting decorated strings still anti-commute. This is because the correction is computed only using stabilizers, so it decorates Z-operators by a linear combination of closed Z-loops, and similarly for X. Moreover, other local decoding algorithms, such as cellular automata and RG decoders, can also be used to generate different LED operators[50]. In the following section, we describe a flexible, "patch-based" local decoder for the toric code, which allows LED to classify a wider range of states as topological.

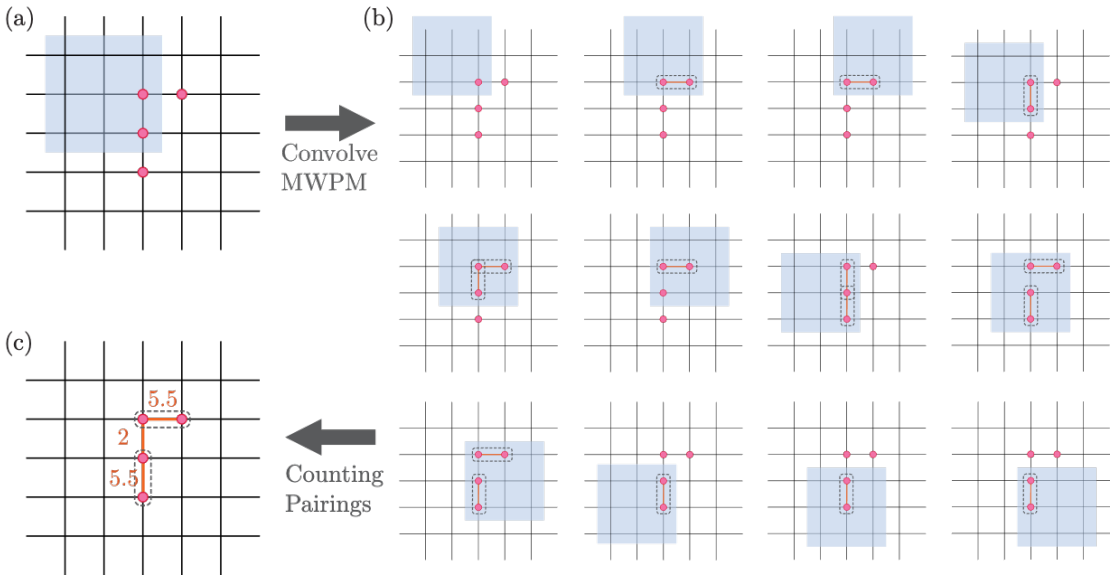


Figure B.3: MWPM-based patch decoder. (a) Example of an error chain which creates four e-anyons. (b) The decoding algorithm performs correction using only local information by splitting the large system into smaller overlapping regions, within each of which the MWPM algorithm is used to find the lowest-weight pairing of anyons. These local regions have open boundaries, hence MWPM can also pair anyons to the boundaries if this is of lower weight. In practice, a slight boundary bias is added to break ties in favor of boundary pairing. (c) The final step requires locally combining the pairing outputs to determine the final pairing. In particular, we count the number of times each site p is paired to sites $q > p$. In the diagram, two equal-weight pairings contribute 0.5 each, though we randomly break the tie in practice. Then, the algorithm pairs p with the q that appears most often. In this example diagram, we connect two pairs which have weight$= 5.5$, and do not form the weight$= 2$ pairing. We see in the simple four-anyon case depicted above, the procedure correctly recovers the pairing with windows of size $l = 3$. In general, this patch-based decoder can correct errors up to distance $d = \lfloor l/2 \rfloor$; moreover, the distance by which it spreads information and the thickness of any associated LED operators are both proportional to l.

B.3 Patch-Based Decoder

The patch-based decoder with variable correction distance d is based on a local minimum weight perfect matching (MWPM) procedure. In the first decoding step, a local MWPM decoder is convolved with all l by l square regions of the toric code, where $l \sim d$; for each region, MWPM takes as input the location of the enclosed anyons. Because both e and m anyons can freely move into and out of the region, this is analogous to decoding a surface code with open boundaries. Therefore, MWPM pairs any given anyon either with another anyon or with the boundary.

The second step aggregates MWPM pairings. Since the square regions can overlap, a pair may appear more than once. As such, after choosing a natural indexing of the plaquettes, we create a list of all MWPM pairings between two plaquettes (p, q) with $p < q$; pairings with the boundary are not included (Figure B.3). For each plaquette p containing an anyon, the patch-based decoder then performs the pairing (p, q) which occurs most often. This procedure naturally favors pairings that flip fewer qubits, because shorter-range pairings can be included in more local patches.

A critical property of this decoder is that it preserves locality. In the first step, MWPM only uses information from local l by l patches. The second step is also local: for each plaquette p, its distance from any possible partner plaquette must be less than l. Aggregation can thus be performed using only the results from a small number of overlapping local patches.

B.4 Decoder Dependence

We now examine how different choices of LED decoders can change the size of the "correctable" region—that is, the region classified as topological. The main text demonstrates results for an $l = 4$ (or $d = 2$) patch-based decoder with coarse-graining size $b = 2$; in Figure B.4, we compare this with the pairing decoder and an $l = 6$ ($d = 3$) patch-based decoder which also uses $b = 2$. We see clearly that the $l = 4$ and $l = 6$ decoders both produce significantly larger correctable regions than the simple pairing decoder, for both coherent and incoherent errors. Meanwhile, the $l = 4$ and $l = 6$ decoders perform similarly, so it appears that the decoder threshold saturates with l. Intriguingly, we observe saturation at an incoherent error rate which is significantly below the known error correction threshold of $p_c \approx 10.9\%$. An interesting open question is to determine whether this discrepancy arises because the patch-based decoder is suboptimal, or because local decoders have some fundamental limit. In Appendix B.10, we show that a "annulus-based" decoder which applies MWPM in a non-local fashion results in a much larger correctable regime.

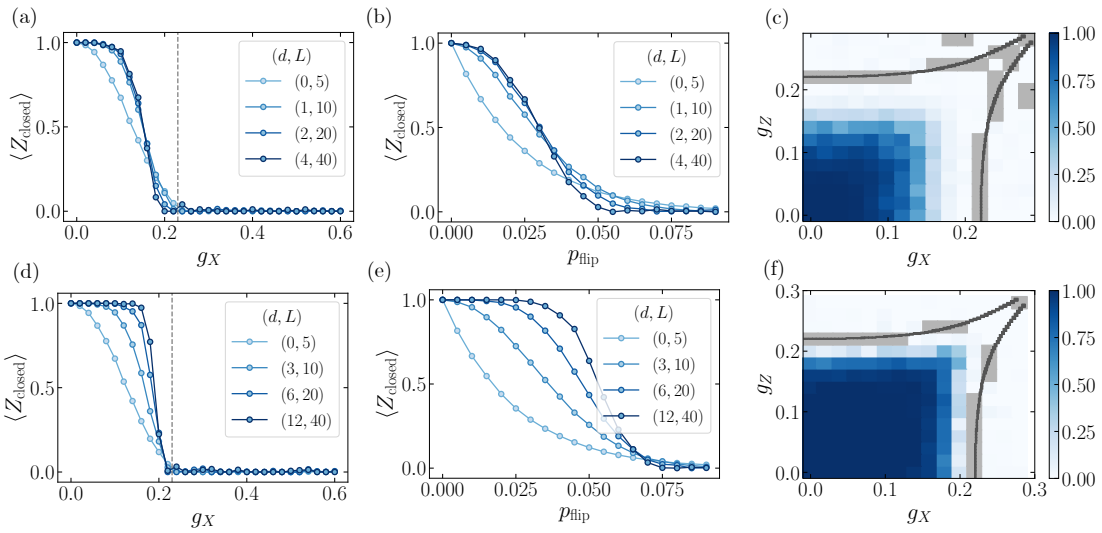


Figure B.4: Decoder choice can significantly affect the extent of the region certified as topological. Here we show LED Wilson loops using coarse-graining with block size $b = 2$ together with (a-c) a $d = 1$ pairing decoder and (d-f) a $d = 3$ patch decoder. Coherent perturbations are considered in panels (a,d), while incoherent errors are studied in (b,e). Panels (c,f) show LED expectation values for various combinations g_X,g_Z after three layers of correction. Note that in the main text, analogous plots are made using the $d = 2$ patch decoder, whose classification boundaries roughly match the $d = 3$ decoder shown here.

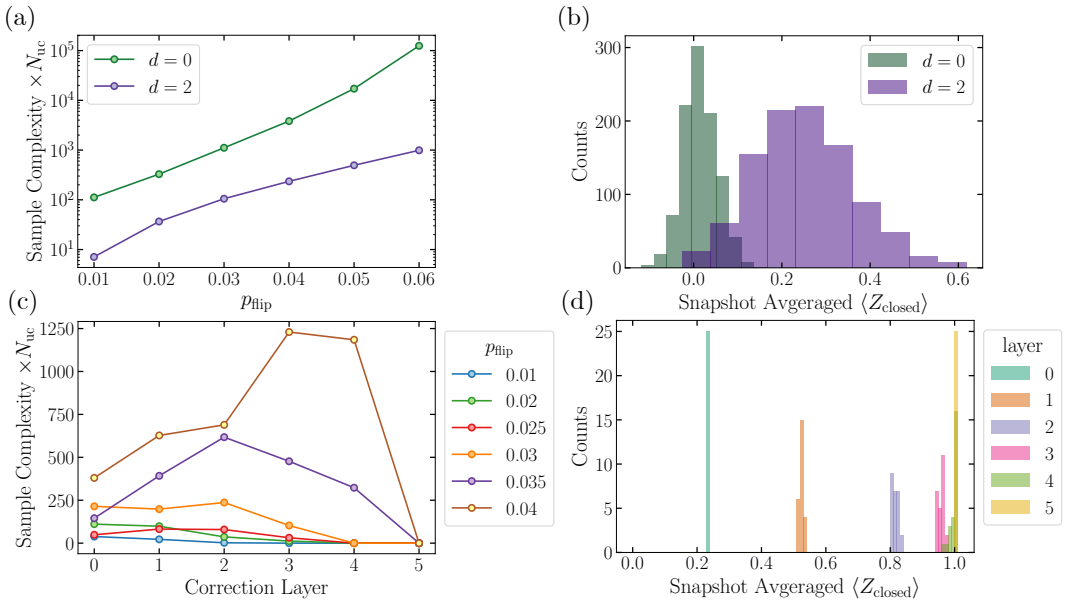


Figure B.5: Sample complexity. (a) Effect of LED on loops of fixed size $L = 10$. Snapshots are drawn from a toric code defined on a 25×25 square lattice with various bit-flip error rates. Sample complexity is computed as described in the main text. Note that sample complexities less than one imply that only a fraction of the system needs to be measured to certify topological order with 95% confidence. As such, the y-axis is also multiplied by the number of unit cells N_{uc}. (b) The histogram of expectation values averaged over a single snapshot confirms that these expectation values are approximately Gaussian distributed. The distribution for LED loops (purple) has much lower weight at zero than uncorrected loops (green), highlighting how fewer samples are needed to verify with high confidence that the closed loops are non-zero. (c,d) We also study the effect of coarse-graining on sample complexity, where the loop length $L = 5 \times 2^n$ grows exponentially with correction layer n. Here, snapshots are taken from a 1024×1024 square lattice. Initially, the sample complexity increases due to a reduction in the number of independent loops available at higher layers. However, it eventually reduces and approaches zero in the topological phase. (d) This turnaround occurs in the limit $\xi \ll d$, where correction is able to remove almost all errors, and loop expectation values approach one.

B.5 Sample Complexity

A key figure of merit for certification of phases is the sample complexity, defined here as the number of samples required to confirm with 95% confidence that the measured loop operator is non-zero. To compute sample complexity, we approximate the LED Wilson loop expectation value evaluated on a large but finite size system as a Gaussian random variable. In this scenario, the important quantity is the ratio of the standard deviation σ to the mean $\langle Z_{\text{closed}} \rangle$. The estimator of the expectation

value has a standard deviation that decreases as $1/\sqrt{S}$ where S is the number of samples. Thus, to confirm the mean is non-zero to two standard deviations (95% confidence), we require approximately $S = (2\sigma/\langle Z_{\text{closed}}\rangle)^2$ samples.

In Figure B.5, we compare the sample complexity of certifying non-zero Wilson loops using bare and LED observables. Two scenarios are considered. In the first, bare and LED Wilson loops are compared at fixed length-scale. Indeed, the sample complexity decreases by an order of magnitude for a range of incoherent error rates below the correction threshold. In the second scenario, coarse-graining is considered, where larger length-scales are probed at each layer. Here, the sample complexity in fact increases at early layers for moderate error rates before falling dramatically. This is because the variance of the signal initially increases, since coarse-graining reduces the number of loops available for averaging in a fixed size system. However, for sufficient correction layers, LED reliably removes almost all errors. This is the regime where $\xi \ll d$, so the Wilson loops saturate at one and their variance approaches zero (see histograms in Figure B.5).

Nevertheless, the initial increase in sample complexity is not simply due to information being removed. Upon further examination of the scenario without coarse-graining, we find a similar initial increase in sample complexity. We interpret this as correction causing adjacent LED loops to become correlated. Finally, we note that the sample complexity measured in this way only improves in the topological phase. In the uncorrectable, disordered phase, the inverse ratio $\langle Z_{\text{closed}}\rangle/\sigma$ rapidly approaches zero.

B.6 Decoder Details for the Ruby Lattice Spin Liquid

We now explain the decoding procedure for a dimer model where qubits lie on the vertices of the ruby lattice, or equivalently, on the links of a kagome lattice. This dimer model supports a $\mathbb{Z}_2$ spin-liquid phase, whose fixed-point is a resonating valence-bond (RVB) state [160]. This state is in the same

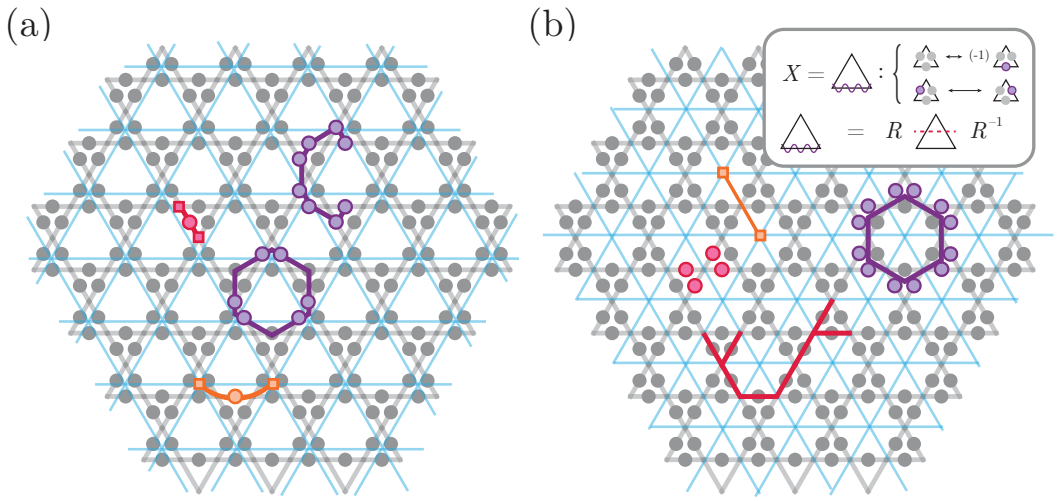


Figure B.6: Decoding for the ruby lattice spin liquid realized in Ref. [143]. (a) For Z-loops, two layers of LED can be performed. In both layers, we use the pairing decoder, which flips a qubit (e.g., red or orange circle) if and only if both neighboring stabilizers (e.g., red or orange squares) are equal to -1. Stabilizers in the first layer (e.g., red squares) are given by $(-1)\prod_{i\in v} Z_i$ for each vertex v of the kagome lattice. The coarse-graining procedure after the first decoding step maps three stabilizers to a single stabilizer (e.g., orange square) in the coarse-grained lattice (blue lines), whose value is determined by the product of the qubits along a loop enclosing a triangle (e.g., purple closed loop). The open strings considered in the main text start and end at hexagons (e.g., purple open string). (b) To measure X-loops, a basis rotation is first performed within each triangle of the kagome lattice, so that the X-string operators become diagonal in the measurement basis (inset and Refs. [143,160]). Each configuration is then mapped to a triangular lattice (blue lines), where each edge of the triangular lattice is determined by the product of four qubits in the original lattice (e.g., red circles); moreover, the X stabilizers of the dimer model become vertex stabilizers in the triangular lattice (e.g., purple hexagons). As before, the pairing decoder flips qubits (orange edges) conditioned on the values of stabilizers (e.g., orange squares). Open strings on the triangular lattice also map to open strings in the kagome lattice (e.g., red string), although the resulting strings are slightly different from the ones measured in Refs. [143,160].

$\mathbb{Z}_2$ universality class as the toric code, as it supports e and m anyons with similar string operators.

We first describe the decoding procedure for e anyons, which correspond to vertices with an even number of adjacent dimers [*]. In the first correction step, we apply the pairing decoder between adjacent vertices. We then coarse-grain the kagome lattice to a triangular lattice by grouping vertices within each upward-pointing triangle. This transforms vertex stabilizers in the kagome lattice to vertex stabilizers in the triangular lattice (Figure B.6a). The pairing decoder is then applied between adjacent triangles in the second correction step. In the main text, we study the flow from uncorrected loops to vertex-paired and triangle-paired loops, which are denoted as as layers 0, 1, and 2, respectively.

We next consider the m anyons, which are associated with hexagonal plaquettes. A rotation is first performed within each triangle, such that the string operators associated with m anyons become diagonal in the measurement basis. This allows us to map each configuration onto triangular lattice, whose vertices are located at the center of each hexagon in the kagome lattice; this mapping transforms the X-stabilizers of the dimer model into vertex Z-stabilizers in the triangular lattice (Figure B.6b). Due to the small experimental system size, we can only perform one layer of correction, and we use the simple pairing decoder on the triangular lattice. We note that open strings on the triangular lattice map onto open strings on the ruby lattice, although the resultant strings are slightly different from the ones measured in [160,143].

B.7 Definition and Properties of Fixed-Point States

As discussed in Chapter 5, one important component for defining any LED procedure is to identify a fixed-point state of the target phase of matter. Here, we review the definition and key properties of a fixed-point state.

[*]Notice that this is an odd Z_2 spin liquid, and the trivial empty state corresponds to maximal occupation of e anyon states.

It is well-known in the literature that gapped quantum ground states can be classified by a real-space RG flow[40,139]. To implement such an RG flow, one notes that any two states in the same phase can be connected by finite-depth local unitary transformations. The hallmark property of topologically-ordered states is the presence of long-range entanglement, and finite-depth local unitary transformations can add or remove local short-range entanglement; thus, for any given state in a topological phase of matter, one can construct a procedure which hierarchically removes all short-range entanglement from this state at increasing length-scales. Then, the resulting state then has zero correlation length, as all short-range entanglement has been removed; this state is known as a fixed-point state of the topological phase.

We also consider important properties of the fixed-point state in the context of our LED procedure and the closely related QCNN procedure of Chapter 4. In both of these procedures, a fixed-point state of the phase under consideration is chosen, and the protocol identifies states within this phase by removing local errors or perturbations on top of this fixed-point state; this is done by performing a decoding operation and a coarse-graining operation, and repeating them n times. In these settings, the fixed-point state has a few special properties: If the input state is equal to the fixed-point state, no errors are detected within each decoding step, and furthermore, the state after each layer of decoding and coarse-graining is equal to the input state, defined on a subsystem of the original system with fewer qubits. Finally, because fixed-point states have zero correlation length, the bare Wilson loops defined in the main text have expectation value equal to $+1$ without performing any correction or LED.

B.8 PEPS Sampling Contraction Details

We start by computing the left boundary MPS of an infinite strip. For sites with $y \leq 0$, we average over measurement outcomes, and hence the doubled PEPS tensor for each site, which contains both

the bra and the ket tensors, is $T^{i',j',k',l'}_{i,j,k,l} = \sum_{pq} (A^*)^{pq}_{i',j',k',l'} A^{pq}_{i,j,k,l}$. The boundary conditions we choose are $|+\rangle_i |+\rangle_{i'}$ at the lower $x = 1$ boundary and $\delta_{i,i'}$ at the $x = L_x$ boundary. Contracting the doubled tensor along an entire column $x = 1, ..., L_x$ results in a matrix product operator (MPO) T_y acting on the boundary MPS. Then, we can compute the effect of an infinite environment by repeatedly applying T_y to some initial boundary MPS until it converges. The right boundary MPS can be similarly computed, by exchanging the input and output directions of the MPO. Note that at each application of T_y, we use SVD truncation to prevent the bond dimension of the boundary MPS from growing too rapidly, rendering the method approximate. However, only singular values smaller than $< 10^{-8}$ are discarded, so truncation errors should be insignificant. We refer to the resulting tensors as the left and right fixed-point of T_y.

The algorithm continues by using the left and right boundary MPS to sample a column of sites. Computing the marginal probability $P(ab)$ on the first site requires contracting a 1D tensor network (see Figure B.1), where the doubled tensor at $x, y = 1, 1$ is replaced by a measurement-dependent one $T(ab)^{i',j',k',l'}_{i,j,k,l} = (A^*)^{ab}_{i',j',k',l'} A^{ab}_{i,j,k,l}$, while the doubled tensor at sites $x > 1$ remain measurement-independent. Note that the tensor network outputs an unnormalized probability distribution; however, since there are only four states per site, the normalization can be computed with little overhead. After drawing a sample ab_{11} for the first site, we replace the doubled tensor at $x, y = 1, 1$ by $T(ab_{11})$, and then compute the distribution of the second site. The process is repeated until the entire column is sampled. To minimize repeat 1D contractions, upper and lower environment tensors can be stored and updated during the sweep $x = 1, ..., L_x$.

Next, the history of samples along the column are used to construct a measurement-dependent MPO $T_y(ab_1, ab_2, ..., ab_{Lx})$, where the doubled tensor at each site is replaced by a measurement-dependent one. Then, the left boundary MPS can be updated by contraction with $T_y(\{ab_i\})$, and the process repeated for the second column. Thus, the left boundary MPS keeps track of the effect of past measurements on future measurements, as the algorithm sweeps from left to right. Mean-

while, the right boundary MPS remains unchanged, as it is modelling a static, infinite environment.

B.9 Calculation of Phase Diagram

To compute the phase diagram, we use the PEPS tensors to construct a transfer matrix with periodic boundary conditions, on small cylinders with circumference L_y measured in unit cells. The largest few eigenvalues of the transfer matrix can be efficiently computed using Krylov algorithms in this regime, and the degeneracy of the largest eigenvalue serves as an alternative signature of the topological transition[82,51]. In particular, the local $\mathbb{Z}_2$ gauge symmetry of the PEPS tensor becomes a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry of the doubled tensor (bra and ket). Hence the topological, e-condensate (Z-paramagnet), and m-condensate (X-paramagnet) correspond to three distinct symmetry broken phases from the point of view of the virtual legs, with degeneracy two, one, and four respectively.

As such, along the transition from topological to e-condensate, which occurs for g_X small and $g_Z \approx 0.2 - 0.3$, the relevant ratio is between the first and second eigenvalues, $\Delta_{12}(g_Z, g_X) = \log(\lambda_1/\lambda_2)$. In contrast, for the transition from topological to m-paramagnet, the relevant gap is between the first and third eigenvalues $\Delta_{13}(g_Z, g_X)$. Furthermore, the model is self-dual, so the wavefunction at $|\psi(g_z, g_x)\rangle$ is equivalent to the wavefunction at $|\psi(g_x, g_z)\rangle$ by a basis rotation and spatial translation. We will use this duality to compute the phase boundary in a way which minimizes finite-size effects, by computing the two gaps at their dual points (Figure B.7a). Therefore, we introduce two parameters g_0, g_1, and let g_1 be the parameter which changes across the transition. The relevant parameter is therefore $\Delta_{12}(g_1, g_0)$ for the topological to e-condensate transition, and $\Delta_{13}(g_0, g_1)$ for the topological to m-condensate transition.

As Δ_{12} grows with increasing g_1, while Δ_{13} reduces with increasing g_1, these two ratios will eventually cross. Indeed, in the limit $L_y \to \infty$, the crossing point should exactly correspond to the phase boundary. In general, finite L_y may shift the boundary. Empirically, we see that at $g_0 = 0$, the

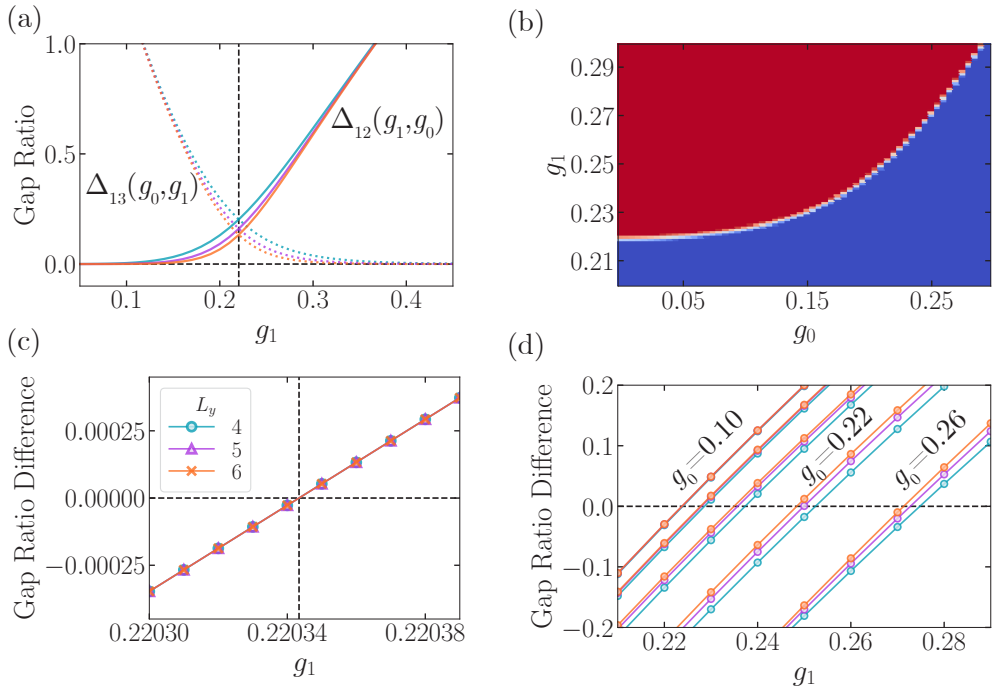


Figure B.7: (a) For a given point (g_0, g_1) on the phase diagram, the phase transition from topological to e-condensate is signaled by the transfer matrix gap ratio $\Delta_{12}(g_0, g_1)$, while the transition from topological to m-condensate is signaled by $\Delta_{13}(g_1, g_0)$, where both quantities are evaluated at the self-dual point with $g_z \leftrightarrow g_x$ switched. When these two gaps are equal, the system must be at the critical point, due to the self-duality. We observe that the crossing point coincides exactly with the known critical point g_c (vertical dotted line), for $L_y = 4, 5, 6$ (blue, purple, orange resp.). (b) Quatitatively, we extract the phase boundary using $L_y = 4$ by finding points where the difference $|\Delta_{12}(g_0, g_1) - \Delta_{13}(g_1, g_0)| \leq 0.01$ is close to zero. (c) Along the cut of the phase diagram with g_0 small, the phase boundary computed from the difference has no visible finite size effects. (d) For larger $g_0 \geq 0.1$, small finite size effects start to appear. Shown are $g_0 = 0.1, 0.14, 0.18, 0.22, 0.26$ from left to right. The phase boundary clearly shifts to larger g_1 with increasing g_0, but $L_y = 4$ (blue) appears to slightly overestimate the transition point compared to $L_y = 6$ (orange).

solvable point, there are essentially no finite size effects, and the agreement with the analytical value $g_c = 0.2203434$ is almost exact (Figure B.7ac). For larger g_0, the dependence on L_y appears minimal until around $g_0 = 0.01$: indeed, even in this case, $L_y = 4$ only overestimates the phase boundary compared to $L_y = 6$ by a few percent (Figure B.7d). As such, we use $L_y = 4$ and compute the phase boundary by identifying points where the difference $\Delta_{12} - \Delta_{13}$ is close to zero (Figure B.7b).

B.10 Annulus Decoders

In this section we consider an alternative approach to constructing error corrected operators, by applying MWPM to an annulus. Anyons supported on the annulus can be paired either with other anyons, or with the boundary of the annulus. Interestingly, MWPM is a global decoder, where introducing or removing anyons can change the pairing far away. However, by construction, the decoder cannot connect the interior of the annulus to the exterior.

We show in Figure B.8 the behavior of these corrected loop operators for coherent, incoherent, and mixed errors. In particular, for the incoherent error model, we see the region classified as topological seems to coincide with the known error recovery threshold of $\approx 10.1\%$ for MWPM. This is much closer to the theoretical optimal error correction threshold of $\approx 10.9\%$, where we expect the phase transition from topological to disordered to occur.

Since this decoder is non-local, our existing theoretical arguments, and the connection to RG, may not apply. Nevertheless, an annulus decoder cannot change the super-selection sector—whether or not an anyon is contained within the annulus. Therefore, if the annulus-correction decorated loops go to one, then asymptotically large regions have well-defined super-selection sectors, suggesting the state is topologically-ordered. Indeed, it has been argued that such an annulus decoder could be considered a witness for topological entanglement entropy [172].

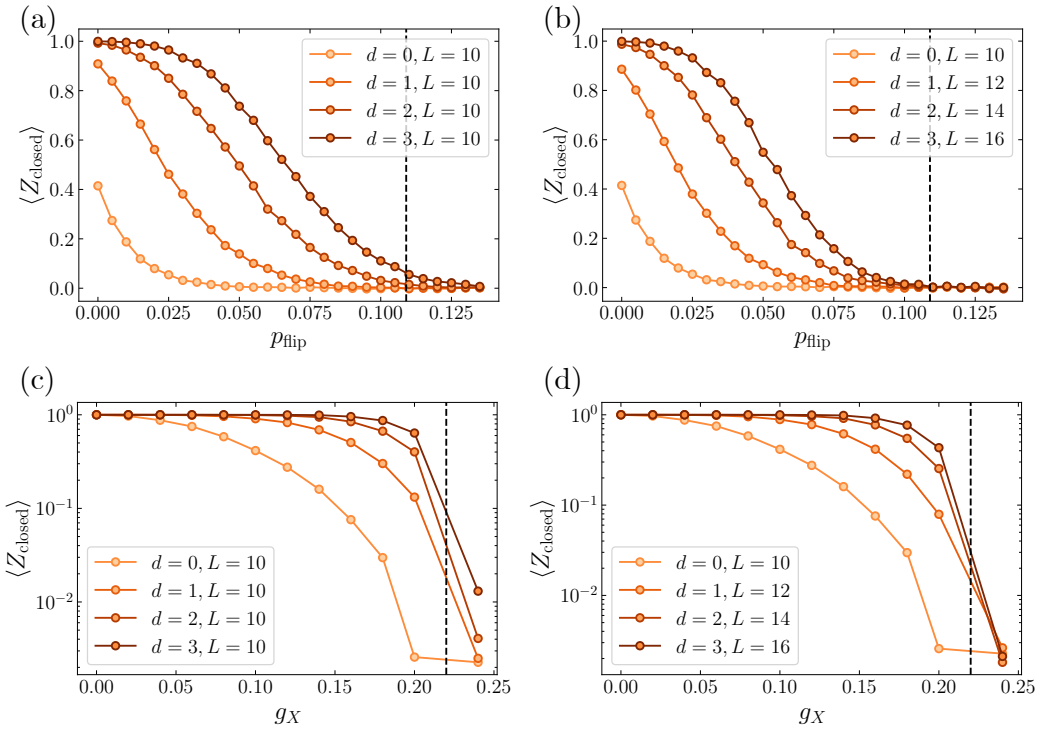


Figure B.8: LED loops using MWPM on an annulus. We see the region classified as topological extends significantly further than the LED loops, for both incoherent (a,b) and coherent (c,d) perturbations. However, additional subtleties arise. When we study fixed loop size and increasing annulus thickness (a,c), there appears to be amplification even on the trivial and disordered sides. However, if we scale the loopsize such that the size of the interior is fixed (b,d), then the improper amplification disappears. An interesting direction for future work is to develop a rigorous understanding of whether non-local annulus-based decorating can also serve as a topological order witness.

C

Supplementary Information for Chapter 6

C.1 Numerical Computation of Branching Ratios and Transition Rates

In this section, we present the results of numerical computation of branching ratios for BBR and RD transitions out of the stretched Rydberg state $70S_{1/2}$, $m_J = 1/2$, $m_I = 3/2$ for ^{87}Rb.

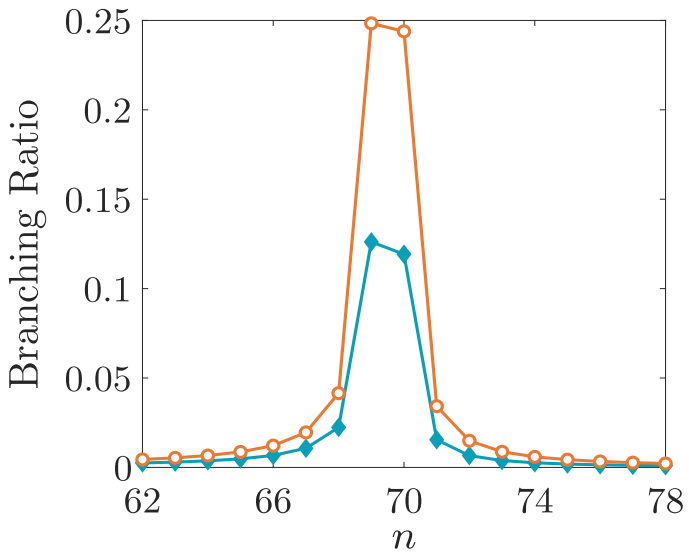


Figure C.1: Branching ratios for BBR transitions between Rydberg states of ^{87}Rb, from the stretched $70S_{1/2}$ state with $m_J = 1/2, m_I = 3/2$ to different P states with $m_J = 3/2$ (empty orange circles) or $m_J = 1/2$ (filled blue diamonds).

C.1.1 Blackbody radiation-induced transitions

To quantify the relative probability of transitioning into different nearby Rydberg P states, we compute the rate $W(nL \to n'L')$ of BBR transitions from a given Rydberg state nL to other Rydberg states $n'L'$ using the Planck distribution of photons at the given temperature T and the Einstein coefficient for the corresponding transition:

$$W(nL \to n'L') = A(nL \to n'L')\bar{n}_\omega \tag{C.1}$$

where $\omega = E_{nL} - E_{n'L'}$ is the transition frequency (E_{nL} and $E_{n'L'}$ are energies of the initial and final states) and

$$A(nL \to n'L') = \frac{4\omega^3}{3c^3}\frac{L_{\max}}{2L+1}R^2(nL \to n'L'). \tag{C.2}$$

In the above equations, $\hbar = 1$, $L_{\max} = \max(L, L')$, and $R(nL \to n'L')$ is the radial matrix element for the electric dipole transition $nL \to n'L'$.

For this work, we used analytic formulas from Refs.[87] and[86] to numerically compute the radial dipole matrix elements for single-photon BBR transitions from the stretched Rydberg state $70S_{1/2}$, $m_J = 1/2$, $m_I = 3/2$ of ^{87}Rb. We then computed the corresponding transition rates using Eq. (C.1), and normalized these by the total BBR rate Γ_{BBR} (see Eq. (6.3)) to obtain the branching ratios. The branching ratios for P states with $m_J = 3/2$ and $m_J = 1/2$ are plotted in Figure C.1 as empty orange circles and filled blue diamonds, respectively. Indeed, we find that the atom decays primarily to the $69P$ and $70P$ states as illustrated in Figure 6.2.

C.1.2 Radiative decay

As shown in Figure 6.2 of the Chapter 6, the radiative decay transitions from the stretched $70S_{1/2}$, $m_J = 1/2$, $m_I = 3/2$ Rydberg state of ^{87}Rb are almost entirely two- or four-photon decay processes to one of the five states in the ground state manifold; this fact was important for converting all Rydberg errors to Z type for fault-tolerant quantum computation. To justify this, we numerically computed the branching ratios for multi-photon spontaneous emission processes by evaluating the ratios of individual transition rates for each decay channel, which are given by the Einstein A coefficients of Eq. (C.2). Due to the cubic dependence of these coefficients on transition frequency, the primary contributions arise from dipole-allowed transitions to states near the ground state manifold. The dipole matrix elements for such transitions scale with the effective principal quantum number n_{eff} of the Rydberg state as $\sim 1/n_{\text{eff}}^{1.5}$. The total RD rate is then given by a sum over Einstein coefficients for all possible target states:

$$\frac{1}{\tau_0} = \Gamma_0 = \sum_{n'L' : E_{nL} > E_{n'L'}} A(nL \to n'L'). \tag{C.3}$$

F	m_F	Branching ratio
2	2	0.534
2	1	0.177
2	0	0.055
2	-1	0.003
2	-2	0.001
1	1	0.168
1	0	0.059
1	-1	0.003

Table C.1: Branching ratios for transition to each ground state of ^{87}Rb for radiative decay processes from the $70S_{1/2}$, $m_J = 1/2, m_I = 3/2$ stretched Rydberg state, accounting for transitions involving up to four-photon emission processes. The contribution from transitions of even higher order is less than 2.5×10^{-4}.

By computing the radial dipole matrix elements using analytic formulas from Refs.[87] and[86], we evaluated the branching ratios for RD processes out of the $70S_{1/2}$, $m_J = 1/2$, $m_I = 3/2$ stretched Rydberg state for ^{87}Rb.

The results of this computation are shown in Table C.1. Indeed, we find that the branching ratios for the remaining three states are each on the order of 10^{-3}, significantly smaller than those for the dominant five transitions. If the total error probability is already very small, these three processes (in particular, the decay to the stretched state with minimal $m_F = -2$) are highly unlikely.

C.2 An Example of Master Equation Solution for Radiative Decay

In Section 6.3.2, we argued that the Kraus operators corresponding to spontaneous emission events from the Rydberg state $|r\rangle$ to the qubit $|1\rangle$ are

$$\begin{aligned} M_0 &= |r\rangle\langle r| + \alpha|1\rangle\langle 1| + \beta|0\rangle\langle 0|, \\ M_r \propto |r\rangle\langle 1|, \quad M_1 &\propto |1\rangle\langle 1|, \quad M_2 \propto |0\rangle\langle 0|, \end{aligned} \tag{C.4}$$

where α, β, and the proportionality constants depend on the specific Rydberg pulse being performed and the probability for an atom in the Rydberg state to decay to the $|1\rangle$ state. We now proceed to derive these constants for the special case of a 2π pulse on the Rydberg transition $|1\rangle \leftrightarrow |r\rangle$ by analytically solving the quantum master equation. For this example calculation, we will ignore BBR transitions and RD transitions to other hyperfine states; these can be included as a straightforward extension.

The master equation for this driven three-level system is (setting $\hbar = 1$)

$$\frac{d\hat{\rho}}{dt} = -i[\hat{H}_d, \hat{\rho}] - \frac{\gamma}{2}\left(\hat{c}^\dagger\hat{c}\hat{\rho} + \hat{\rho}\hat{c}^\dagger\hat{c} - 2\hat{c}\hat{\rho}\hat{c}^\dagger\right), \tag{C.5}$$

where $\hat{\rho}$ denotes the density matrix of the system, $\hat{H}_d = i\Omega\left(|r\rangle\langle 1| - |1\rangle\langle r|\right)$ is the driving Hamiltonian, $\hat{c} = |1\rangle\langle r|$ is the quantum jump operator corresponding to spontaneous emission $|r\rangle \mapsto |1\rangle$, and γ is the probability for an atom in the Rydberg state to decay to $|1\rangle$. We assume the qubit is initially encoded in the hyperfine manifold Span$\{|0\rangle, |1\rangle\}$, so that the initial density matrix can be written as

$$\hat{\rho}_0 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \rho_{11} & \rho_{10} \\ 0 & \rho_{01} & \rho_{00} \end{bmatrix} \tag{C.6}$$

(we order the matrix columns and rows as $\{|r\rangle, |1\rangle, |0\rangle\}$). Upon solving the resulting coupled first-order differential equations, we find that the final state after the 2π pulse with decay is, to leading order in γ/Ω,

$$\hat{\rho}_f = \begin{bmatrix} 3\gamma t_\pi \rho_{11}/4 & 0 & 0 \\ 0 & (1 - 3\gamma t_\pi/4)\rho_{11} & -e^{-\gamma t_\pi/2}\rho_{10} \\ 0 & -e^{-\gamma t_\pi/2}\rho_{01} & \rho_{00} \end{bmatrix}. \tag{C.7}$$

Here, $t_\pi = \frac{\pi}{2\Omega}$ is the duration of a π pulse. Indeed, Eq. (C.7) confirms our intuition from Section

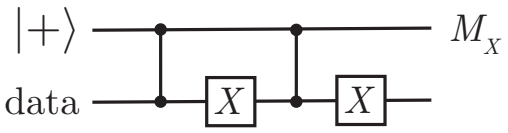


Figure C.2: Circuit for detecting atom loss.

6.3.2 that the coherences ρ_{r1}, ρ_{1r} vanish upon "averaging" over all possible transition times during the 2π pulse.

One can then verify that, to leading order in γ/Ω, the Kraus operators

$$M_0 = |r\rangle\langle r| + \sqrt{1-p_2}|0\rangle\langle 0| + \sqrt{(1-p_1)(1-p_2)}|1\rangle\langle 1| \tag{C.8}$$

$$M_r = \sqrt{p_1(1-p_2)}|r\rangle\langle 1| \tag{C.9}$$

$$M_1 = \sqrt{p_2}|1\rangle\langle 1| \qquad M_2 = \sqrt{p_2}|0\rangle\langle 0| \tag{C.10}$$

give rise to the desired evolution from ρ_0 to ρ_f provided we take $p_1 = 3\gamma t_\pi/4$ and $p_2 = \gamma t_\pi/8$.

C.3 Atom Loss Errors

As mentioned in Section 6.3.4, neutral atom setups can also suffer from atom loss errors if the trapping is imperfect, or if the trapping lasers need to be turned off during Rydberg excitation (e.g. as done for ^{87}Rb in Ref. [12]). Fortunately, such errors can also be detected and corrected within our FTQC framework at the cost of one ancilla qubit and some extra gates for each operation. In particular, an atom loss event can be detected by applying the circuit of Figure C.2 for each data qubit after using the optical pumping technique to correct for leakage out of the computational subspace. The ancilla measurement will then produce $+1$ in the presence of atom loss, and -1 if such an er-

ror did not occur. Once detected, an atom loss error can be converted to a single-qubit Pauli-Z or X type error if a reservoir of atoms is available, for instance by replacing the lost atom with a new atom initialized to the $|0\rangle$ state.

We now discuss the steps needed for establishing robustness against errors occurring during this circuit. As in the case of fault-tolerant Rydberg leakage detection discussed in Appendix C.4, to protect against ancilla errors in Figure C.2, we again adopt a multi-step ancilla measurement protocol and require two positive ancilla measurements to confirm an atom loss error. On the other hand, any phase-flip error on the data qubit cannot propagate to more than a single physical qubit error per logical qubit in the universal gate set implementations for Ryd-7 or Ryd-3. Leakage errors (Rydberg or non-Rydberg) can be addressed by repeating the respective re-pumping procedures after applying the atom loss detection circuit. Thus, by incorporating this circuit into the implementation of fault-tolerant stabilizer measurements and logical operations in Sections 6.4 and 6.5, we can also address atom loss errors in our FTQC protocols.

Notice that this circuit can be used for atom loss after correcting for leakage into atomic states outside the computational subspace by using the blockade effect and optical pumping techniques. In addition, this approach does not distinguish between atom loss and leakage into other hyperfine states, so it can also be used to suppress any residual hyperfine leakage errors.

C.4 Converting Rydberg Leakage to Pauli Errors

Once a Rydberg leakage error is detected, it can be converted to an atom loss error by ejecting the Rydberg atom, which is naturally done by the anti-trapping potential from the tweezer [16] and can be expedited by pulsing a weak, ionizing electric field [10,45,173]. The exact location of the ejected atom can be determined by following the atom loss protocol outlined in Appendix C.3 and Figure C.2 [*];

[*] In this case, the atom loss protocol does not need to be applied in a robust fashion, since an error has already occurred.

subsequently, the ejected atom can be replaced with a fresh atom prepared in the $|1\rangle$ state. Although this process simply replaces the Rydberg atom by an atom in the $|1\rangle$ state, by using the operator identity $|1\rangle\langle 1| = \frac{1}{2}(1 - Z)$, it is straightforward to see that the resulting state is now a superposition of the original state without error, and the same state with a Z error on this physical qubit. Such Z-type errors can be detected and corrected for using stabilizer measurements in both the seven-qubit and three-qubit codes. This procedure can also be modified to convert the Rydberg leakage error to a Pauli X-type error by applying Hadamard gates at the beginning and end; this is used to in the logical CCZ gate for Ryd-7 (see Algorithm 3).

To reduce the need for applying the atom loss correction circuit, one could add a preventative step after every entangling gate which incoherently re-pumps any remnant population in several most probable Rydberg states into the qubit $|1\rangle$ state. This re-pumping can be implemented via the following three-step procedure:

1. Swap the population in $|1\rangle$ and the stretched ground state $|F = I + 1/2, m_F = I + 1/2\rangle$.

2. For the most probable final states $|r'\rangle$ of a BBR transition (or the Rydberg state $|r\rangle$ in the case of RD), perform a Rydberg laser pulse that sends $|r'\rangle$ (or $|r\rangle$) to a short-lived P state. In particular, we choose the P state with the smallest possible n, largest possible F, and largest possible m_F. This state will quickly decay to the stretched state $|F = I+1/2, m_F = I+1/2\rangle$, and cannot decay to any other ground state.

3. Repeat Step (1).

By applying this procedure preventatively, one can convert a large fraction of Rydberg leakage errors to Z-type errors without the need for the atom loss correction circuit of Figure C.2.

C.5 Fault-Tolerant Detection of Rydberg Leakage Errors

As mentioned in Section 6.4, for fault-tolerant error detection and correction, it is important to address any errors that may occur on an ancilla used to probe for Rydberg population. This can be done by using a multi-step measurement procedure to detect leakage for the ancilla qubit:

1. Perform a Hadamard gate on the ancilla.
2. Check whether the ancilla is in the $|1\rangle$ state (e.g. by coupling $|1\rangle$ to a cycling transition and detecting fluorescence).
3. Perform an X gate on the ancilla.
4. Check for $|1\rangle$ population again.

If neither the second nor the last step yields $|1\rangle$, the ancilla atom must have undergone a leakage error. In that case, we convert any possible ancilla atom Rydberg error to a possible Z-type error as described in Appendix C.4. Similarly, because the Rydberg pulses can potentially cause a phase-flip error on the ancilla qubit, if a Rydberg leakage error is detected by the ancilla, the detection protocol must be repeated once more to ensure that the outcome did not result from such an error.

C.6 Error Syndromes with Postponed Measurements

In Section 6.4.1, we discussed how Rydberg leakage detection can be postponed in the Ryd-7 stabilizer measurement and controlled-phase gate protocols to facilitate experimental implementation. This relied on the ability to use stabilizer measurements to distinguish between the possible correlated errors that can result from postponed detection of a Rydberg leakage error. Here, we present details on how to use error syndromes to identify the corresponding correlated error in each case. As

Error	g_4	g_5	g_6
$X_5X_6X_7$	-1	$+1$	$+1$
X_6X_7	$+1$	$+1$	-1
X_7	-1	-1	-1

Table C.2: Using error syndromes to distinguish between correlated errors resulting from postponed detection of Rydberg leakage during measurement of the $X_4X_5X_6X_7$ stabilizer in the Ryd-7 FTQC protocol. Because the possible correlated errors are all products of Pauli-X errors, we show the corresponding values of $Z^{\otimes 4}$ stabilizer measurements.

in Section 6.4, we assume the stabilizers for the Steane code are ordered as

$$\begin{aligned} g_1 &= IIIXXXX & g_2 &= IXXIIXX & g_3 &= XIXIXIX \\ g_4 &= IIIZZZZ & g_5 &= IZZIIZZ & g_6 &= ZIZIZIZ. \end{aligned} \tag{C.11}$$

For the stabilizer measurement, we will consider (without loss of generality) the measurement of g_1 on qubits $4, 5, 6, 7$ using a circuit of the form shown in Figure 6.1b. If a Rydberg leakage error occurs on the ancilla atom at any point, the data atoms do not suffer any correlated errors. On the other hand, if a data atom suffers a Rydberg leakage error during the circuit, the possible correlated errors that can result are $X_5X_6X_7$, X_6X_7, or X_7. These errors can be distinguished by measuring the $Z^{\otimes 4}$ stabilizers of the seven-qubit code; the corresponding error syndromes are shown in Table C.2.

For the case of the logical CCZ gate, we grouped the 27 physical Rydberg gates into groups $\mathcal{G}_i$ of three, and performed Rydberg leakage detection after each group. Without loss of generality, we will consider the group $\mathcal{G}_1$ in Figure 6.4. There are two possible correlated errors that could result from the delayed detection of Rydberg leakage in this case (up to a single-qubit error within each logical qubit): $R(2_A, 2_B; 2_C)R(3_A, 3_B; 3_C)$ and $R(3_A, 3_B; 3_C)$. By writing the Rydberg gate as

$$R(j, k; l) = (1 + Z_j)(1 + Z_k)(1 + Z_l) - 1, \tag{C.12}$$

we find that the two cases can be distinguished by measuring the stabilizers g_2 and g_3 for each of

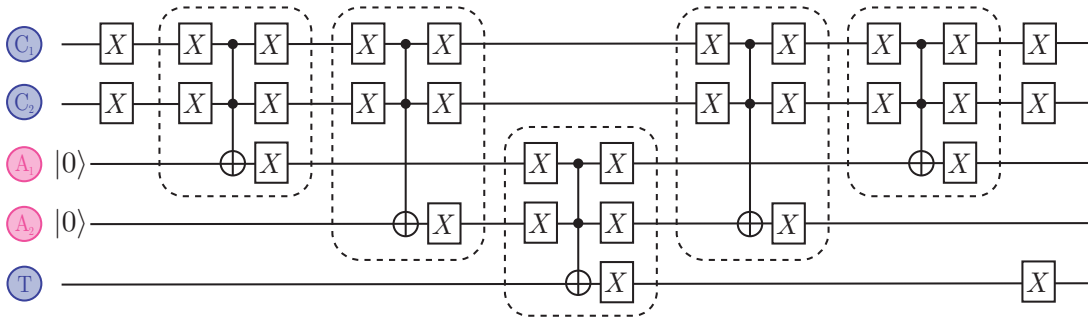


Figure C.3: Using two ancilla qubits and multiple Rydberg states to implement a bias-preserving Toffoli gate between control atoms C_1, C_2, and target atom T. The ancilla atoms (A_1 and A_2) are chosen to lie on either side of the target atom. The dotted boxes indicate the most natural bias-preserving three-qubit gate for Rydberg systems, where π pulses $|1\rangle \leftrightarrow |r_+\rangle$ are applied to each of the first two (i.e., the upper two) involved atoms, the bias-preserving pulse sequence of Figure 6.5 is applied to the third (lower) atom, and $-\pi$ pulses $|1\rangle \leftrightarrow |r_+\rangle$ are applied to the first two qubits; the Rydberg states $|r_\pm\rangle$ are chosen to be either $|r_{1,\pm}\rangle$ or $|r_{2,\pm}\rangle$ for each such gate. In this circuit, we set $|r_\pm\rangle = |r_{1,\pm}\rangle$ in the first, second, fourth, and fifth cases, while choosing $|r_\pm\rangle = |r_{2,\pm}\rangle$ for the third one. With this choice of Rydberg levels, the two ancillas will not interact with each other during the third Rydberg gate. However, we note that the two control atoms may interact with each other during the first, second, fourth, and fifth entangling gates if the distance between them is less than one blockade radius; this is not problematic because Rydberg errors can occur during at most one of these gates, so at least one ancilla atom will generate the correct interaction with the target atom during the third gate.

the logical qubits. In the former case, at least one of the the logical qubits would have either a Z_2 or Z_2Z_3 error, giving rise to stabilizer eigenvalues $(g_2, g_3) = (-1, +1)$ or $(+1, -1)$, while in the latter scenario, all three sets of stabilizer measurements would yield $(-1, -1)$ or $(+1, +1)$.

C.7 Implementation of a Bias-Preserving Toffoli Gate

In Figure 6.6, we showed how an ancilla atom can be used to eliminate X-type errors resulting from control atom decay in the implementation of a bias-preserving CNOT gate. Analogously, a bias-preserving Toffoli gate can be implemented by making use of two ancilla atoms which lie on either side of the target atom. This protocol is illustrated in Figure C.3.

As with the case of the bias-preserving CNOT, the choice of Rydberg states differs throughout the procedure. By coupling the atoms to $|r_{2,\pm}\rangle$ during the third gate of Figure C.3 and using an-

cilla atoms on opposite sides of the target atom, we ensure that the ancilla atoms do not interact with each other via Rydberg blockade during this gate; this is important in case one of the ancilla atoms undergoes a radiative decay transition during this gate. On the other hand, the other entangling gates in Figure C.3 all use the Rydberg states $|r_{1,\pm}\rangle$, due to larger distances between the atoms during these gates. We note that the two control atoms may interact with each other during these four gates if the distance between them is less than one blockade radius, which is different from the case of the third gate. This is acceptable because Rydberg errors can occur during at most one of these four gates, so at least one ancilla atom will generate the correct interaction with the target atom during the third gate.

C.8 Computing Resource Costs for Rydberg FTQC Protocols

We now provide details on how to obtain the resource costs for Ryd-7 and Ryd-3 presented in Tables 6.1 and 6.2.

For the Ryd-7 protocol, each stabilizer measurement requires four two-qubit Rydberg gates in the absence of errors (see Algorithm 1); thus, 24 two-qubit gates are required to measure all stabilizers. If an error occurs, the worst case scenario for the stabilizer measurement is when the first five stabilizers all have $+1$ eigenvalues, while the very last stabilizer is measured to be -1. In this case, g_4, g_5, and g_6 need to be re-measured, which requires 12 additional two-qubit gates. The logical CCZ gate for Ryd-7 is implemented using 27 physical three-qubit gates in the absence of error, as described in Algorithm 3. The worst case error in this case is a Rydberg leakage error that occurred during the first entangling gate in the final group $\mathcal{G}_9$ of Figure 6.4. In this scenario, identifying the location of the Rydberg leakage error requires up to 18 additional two-qubit gates, while measuring the stabilizers $g_2, g_3, \ldots, g_6$ for all three logical qubits would amount to 60 additional two-qubit gates; the correction circuit could require up to two additional three-qubit gates.

In the Ryd-3 protocol, each of the two stabilizer measurements requires two bias-preserving CNOT gates (Figure 6.1d), and each bias-preserving CNOT gate is broken down to two two-atom gates and one three-atom entangling gate (see Section 6.5.1). Thus, in the absence of error, the stabilizer measurements would require eight two-qubit gates and four three-qubit gates. If an error occurs, the worst case scenario is if the second stabilizer is measured to be -1; in this case, both stabilizers need to be re-measured, and the gate cost is doubled. The Ryd-3 CCZ gate can be implemented in a round-robin fashion in the same way as the Ryd-7 CCZ, which is bias-preserving and uses 27 physical three-qubit gates.

Finally, the Ryd-3 Hadamard gate consists of a fault-tolerant, bias-preserving Toffoli gate followed by single-qubit measurements and rotations (Figure 6.8). The pieceable fault-tolerant Toffoli gate in the Ryd-3 code consists of nine physical bias-preserving Toffoli gates and two rounds of error correction. As discussed above, each round of error correction involves eight two-atom Rydberg gates and four three-atom Rydberg gates. When the data atoms within each logical qubit are indexed as in Figure C.4b and we are implementing a logical Toffoli gate CCX_{ABC} between the three qubits A, B, C highlighted in bold, the number of Rydberg gates required to implement each physical Toffoli gate depends on the blockade radius $R_{B,1}$. If the blockade radius $R_{B,1}$ is larger than $3.61d$, each physical Toffoli gate can be implemented using two ancilla atoms (one on either side of the target atom) and five three-atom Rydberg gates, as described in Appendix C.7; this is because the distance between any physical control atom C_i and any ancilla A_j in Figure C.3 will always be less than the blockade radius $R_{B,1}$, so the entangling gates can be implemented directly. In this case, each physical Toffoli gate involves five three-atom Rydberg gates, so the total gate count (upon including the QEC steps) is 16 two-atom gates and 53 three-atom gates in the absence of errors. On the other hand, if we wish to reduce the blockade radius requirement to $R_{B,1} > 3d$, there are two physical Toffoli gates (corresponding to the choices $j_A = l_C = 1, k_B = 2$ and $j_A = l_C = 3, k_B = 2$), where the distance between one of the physical control atoms and one of the ancilla atoms (2_B and

A_3 in Figure C.4b) would be too large to directly implement a Rydberg entangling gate required for the physical Toffoli gate. Instead, in place of the first (respectively, second) three-atom Rydberg gate involving A_3, we would implement a Rydberg gate with the same two control atoms and one of the ancilla atoms A_1 or A_2, whichever is not involved in the rest of the Figure C.3 circuit, followed (respectively, preceded) by a bias-preserving CNOT gate between that ancilla and A_3. These gates can be implemented directly because both A_1 and A_2 are within the blockade radius of 2_B, 1_A, 2_A, 3_A, and A_3. In this way, four extra two-atom gates are required for the logical Toffoli (two for the physical Toffoli with $j_A = l_C = 1, k_B = 2$ and two for the physical Toffoli with $j_A = l_C = 3$, $k_B = 2$), which increases the total gate count to 20 two-atom gates and 53 three-atom gates in the absence of error, as shown in Table 6.2. With errors, the worst case scenario is if the final stabilizer measurement in the second round of QEC yields -1, in which case the stabilizers need to be measured again; this adds another eight two-atom gates and four three-atom gates to the total resource cost.

C.9 Computing Rydberg Blockade Radius Requirements for Rydberg FTQC Protocols

To obtain the blockade radius requirement for the Rydberg FTQC protocols, we must identify each physical qubit with an atom on the lattice, and then determine the maximum distance between two atoms which must interact with each other during a Rydberg gate. When the underlying atoms are placed in a triangular lattice, Figure C.4 depicts convenient identifications for both the Ryd-7 and Ryd-3 codes. In this figure, numbers are used indicate the indices of data atoms within each logical qubit. (The index of a physical qubit within each logical qubit is the position, counting from the left, of that qubit in the definition of the logical states; see Equations (6.6) and (6.7) for the seven-qubit code, or Equation (6.13) for the three-qubit code.)

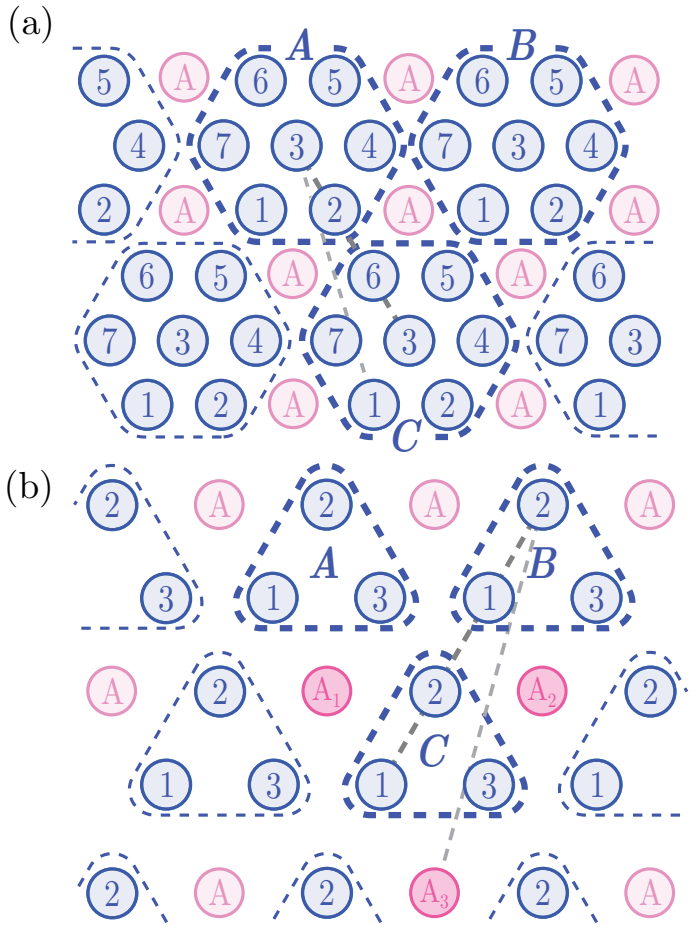


Figure C.4: Example labeling of atoms for the Ryd-7 and Ryd-3 FTQC protocols used to derive the gate counts and blockade radius requirements. As in Figure 6.1, data atoms are shown in blue, while ancilla atoms are shown in pink. (a) In the Ryd-7 protocol, each logical qubit consists of seven data atoms (blue dotted hexagons). For each data atom, a number is used to indicate which physical qubit of the seven-qubit logical state the atom encodes. With this labeling, the blockade radius R_B is defined by the interaction range needed to perform a logical CCZ gate between three neighboring logical qubits such as A, B, and C. Using the specific CCZ protocol given in Algorithm 3, the blockade radius requirement is then $R_B > 3.61d$, where d is the spacing between nearest neighbors on the lattice; this is determined by the distance between physical atoms 3_A and 1_C (thinner, light grey dotted line). However, by using a different set of physical CCZ gates to implement the logical CCZ, this requirement can be reduced to $R_B > 3d$ (thicker, dark grey dotted line). (b) In the Ryd-3 protocol, each logical qubit consists of three data atoms (blue dotted triangles). For each data atom, a number is used to indicate which physical qubit of the three-qubit logical state the atom encodes. With this labeling, the larger blockade radius $R_{B,1}$ is determined by the interaction range required for performing a logical Toffoli gate between three neighboring logical qubits such as A, B, and C. In this case, there are two possibilities for $R_{B,1}$—either $R_{B,1} > 3.61d$ (thinner, light grey dotted line) or $R_{B,1} > 3d$ (thicker, dark grey dotted line). When the larger blockade radius of $3.61d$ can be realized, the resource cost for the logical Toffoli and Hadamard gates can be reduced by four two-qubit entangling gates compared to the numbers presented in Table 6.2 (see also Appendix C.9).

In the Ryd-7 protocol, the blockade radius is defined by the interaction range needed to perform a logical CCZ gate between three neighboring logical qubits such as A, B, and C. Using the specific protocol given in Algorithm 3, which involves 27 physical CCZ gates between atoms $j_A, k_B, l_C \in \{1, 2, 3\}$, we find the largest interaction range is required to perform the physical CCZ gate between farthest-separated triples such as $(j_A, k_B, l_C) = (3, 3, 1)$. For this specific case, the distances between atom pairs are $\text{dist}(j_A, k_B) = 3d$, $\text{dist}(j_A, l_C) = \sqrt{(7/2)^2 + 3/4}d \approx 3.61d$, and $\text{dist}(k_B, l_C) = 4d$. To apply the three-qubit Rydberg gate $R(j_A, k_B; l_C)$ as defined in Section 6.2, this would require a blockade radius of $R_B > 4d$. However, this is not entirely necessary for our purposes: instead, it is sufficient that *two* out of the three distances $\text{dist}(j_A, k_B)$, $\text{dist}(j_A, l_C)$, and $\text{dist}(k_B, l_C)$ be less than the blockade radius. To see this, let us suppose, for example, that the distance between the two control atoms j_A and k_B is greater than R_B. In this case, applying the same pulse sequence as illustrated in Figure 6.3b would result in a three-qubit gate $R = \text{diag}(1, -1, -1, -1, -1, -1, 1, 1)$, which can also be obtained from the CCZ gate by single-qubit unitaries ($R \propto Y_1 Y_2 (\text{CCZ}) X_1 X_2$).

The argument above allows the blockade radius requirement for Ryd-7 to be reduced to $R_B > 3.61d$ (thinner, light grey dotted line in Figure C.4). In fact, by modifying the implementation of the logical CCZ gate, it is possible to further reduce this requirement to $R_B > 3d$ (thicker, dark grey dotted line in Figure C.4); this is shown in Appendix C.10.

In the Ryd-3 protocol, the blockade radius $R_{B,1}$ is determined by the interaction range required to implement the logical Toffoli gate between neighboring logical qubits (e.g., A, B, and C in Figure C.4). As discussed in Appendix C.8, there are two possibilities in this case. To directly implement every physical bias-preserving Toffoli gate using the circuit of Figure C.3, the distance between 2_B and A_3 must be less than $R_{B,1}$; this requires $R_{B,1} > \sqrt{(7/2)^2 + 3/4}d \approx 3.61d$ (thinner, light grey dotted line in Figure C.4). However, this requirement can be reduced to $R_{B,1} > 3d$ (thicker, dark grey dotted line in Figure C.4) at the expense of four additional two-atom entangling gates per logical Toffoli or Hadamard operation.

C.10 Blockade Radius Reduction for Ryd-7

To reduce the blockade radius requirement from $R_B = 3.61d$ to $R_B = 3d$ in the Ryd-7 protocol, we must modify the implementation of the logical CCZ operation. Recall that Algorithm 3 implements a logical CCZ gate using 27 physical CCZ gates between the first three physical qubits of every logical qubit. This round-robin decomposition makes use of Eq. (6.11), which we now derive:

$$\mathrm{CCZ}_{ABC} = \prod_{j_A,k_B,l_C\in\{1,2,3\}} \mathrm{CCZ}(j_A,k_B,l_C), \tag{C.13}$$

To begin the derivation, we first recall that the logical states (6.6) and (6.7) of the seven-qubit code have well-defined parity: the number of physical qubits in the $|1\rangle$ state is always even for $|0\rangle_L$ and odd for $|1\rangle_L$. It then follows that the logical CCZ gate can be implemented in a fully round-robin fashion involving all physical qubits

$$\mathrm{CCZ}_{ABC} = \prod_{j_A,k_B,l_C\in\{1,2,\ldots,7\}} \mathrm{CCZ}(j_A,k_B,l_C). \tag{C.14}$$

This is because the round-robin implementation results in a -1 phase accumulation for each triple (j_A,k_B,l_C) of physical qubits in the $|1\rangle$ state, and the number of such triples is odd if all logical qubits are in the $|1\rangle_L$ logical state, while it is even if at least one logical qubit is in the $|0\rangle_L$ state. To reduce this to Eq. (C.13), we notice that for each choice of j_A and k_B, the product

$$\prod_{l_C\in\{4,5,6,7\}} \mathrm{CCZ}(j_A,k_B,l_C) \tag{C.15}$$

acts as an identity operation on the logical qubits, because $g_4 = Z_4Z_5Z_6Z_7$ is a stabilizer of the seven-qubit code. We then multiply both sides of Eq. (C.14) by this operator, and use the fact that all the CCZ gates commute with each other and square to the identity operator. In this way, the

product over l_C in the logical CCZ gate can be reduced from $l_C \in \{1, 2, ..., 7\}$ to $l_C \in \{1, 2, 3\}$. Because the CCZ gate is symmetric in the three involved qubits, this same argument can be applied to reduce the products over j_A and k_B to obtain Eq. (C.13).

To reduce the blockade radius requirement from $R_B = 3.61d$ to $R_B = 3d$, we can replace the product (C.15) by

$$\prod_{l_C \in \{1,2,4,7\}} \mathrm{CCZ}(j_A, k_B, l_C) \tag{C.16}$$

in our derivation for one of the logical qubits, say qubit C. This is because the single-qubit operator $Z_1 Z_2 Z_4 Z_7 = g_2 g_3$ is the product of two stabilizers, so the operator (C.16) also acts trivially on the logical subspace. It follows that

$$\mathrm{CCZ}_{ABC} = \prod_{\substack{j_A, k_B \in \{1,2,3\} \\ l_C \in \{3,5,6\}}} \mathrm{CCZ}(j_A, k_B, l_C). \tag{C.17}$$

Thus, the 27 physical CCZ gates in Algorithm 3 may be replaced by the 27 CCZ gates used in the right hand side of Equation (C.17).

Given the geometrical layout of individual atoms within each logical qubit shown in Figure C.4a, we see that the required interaction range for implementing the logical CCZ operation using these 27 gates is smaller than the interaction range required to perform the 27 gates of Algorithm 3. Furthermore, following the observation made in Appendix C.9, we notice that these 9 physical qubits need not all be within the blockade radius of each other, so long as every physical qubit $j_A \in \{1, 2, 3\}$ is within distance R_B of every $l_C \in \{3, 5, 6\}$, and every $k_B \in \{1, 2, 3\}$ is within distance R_B of every $l_C \in \{3, 5, 6\}$. This requirement is satisfied for any $R_B > 3d$, as shown in Figure C.4a.

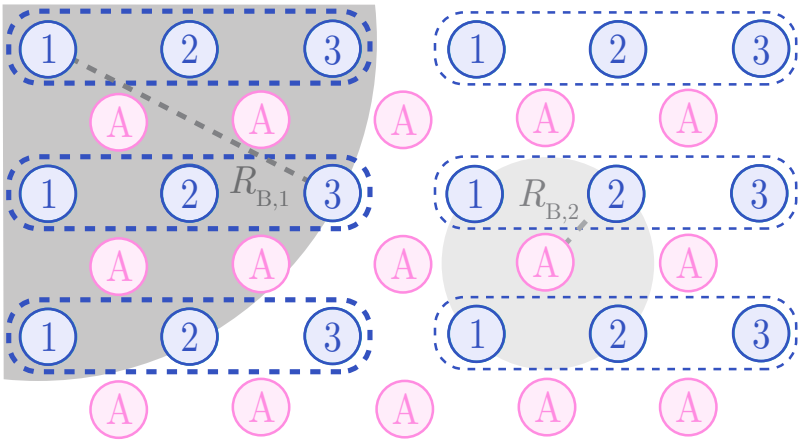


Figure C.5: Square lattice geometry for the Ryd-3 FTQC protocol. Data (blue) and ancilla (pink) atoms are placed on the vertices of a square lattice in an alternating fashion, with three data atoms comprising a logical qubit (blue dotted boxes). The numbers on each data atom indicate the index of that atom within each logical qubit; this is relevant for the implementation of stabilizer measurements and logical operations. Two Rydberg states with different blockade radii are required to implement the bias-preserving CNOT and Toffoli gates. The larger blockade radius $R_{B,1}$ must be larger than $\sqrt{10}d$ (dark grey), where d is the nearest-neighbor spacing on the lattice, while the smaller blockade radius must satisfy $d < R_{B,2} < 2d$ (light grey). The interaction range $R_{B,1}$ is needed to perform a logical CCZ gate between the three logical qubits indicated in bold.

C.11 Square Lattice Geometry for Ryd-3

As mentioned in Section 6.5.3, the Ryd-3 protocol can also be implemented when the underlying physical atoms are placed on a square lattice. In this case, the data and ancilla atoms are placed on the vertices of the lattice in an alternating fashion as shown in Figure C.5. The stabilizer measurements can be implemented as discussed in Section 6.5.2 if the smaller blockade radius $R_{B,2}$ satisfies $d < R_{B,2} < 2d$. The logical operation requiring the largest interaction range is the logical CCZ gate

$$\mathrm{CCZ}_{ABC} = \prod_{j_A,k_B,l_C\in\{1,2,3\}} \mathrm{CCZ}(j_A,k_B,l_C), \tag{C.18}$$

which is implemented from 27 physical CCZ gates. To implement each physical gate, the distance between every pair (j_A,l_C) and (k_B,l_C) must be less than the larger blockade radius $R_{B,1}$, as dis-

cussed in Appendix C.9. The longest such distance is $\sqrt{10}d$ as shown in the dark grey dotted line of Figure C.5, so the corresponding blockade radius requirement for this geometry is $R_{B,1} > \sqrt{10}d$.

With these blockade radii, the protocols of Section 6.5.2 can be directly applied to perform all logical operations. We note that the higher density of ancilla atoms in this arrangement allows us to implement every physical Toffoli gate in the logical Toffoli operation directly using the circuit of Figure C.3, without the need for additional ancilla atoms or CNOT gates (as was the case for two physical Toffoli operations under the triangular lattice geometry). In this way, for the square lattice geometry, the number of two-qubit entangling operations required for the logical Hadamard or Toffoli operations may be reduced by 4 compared to the numbers shown in Table 6.2.

C.12 Optical Pumping Procedure for the Bias-Preserving CNOT

To implement the bias-preserving CNOT pulse sequence shown in Figure 6.5 of Chapter 6, it is important that the optical pumping procedure in the final step pumps only the $m_F > 0$ states to the $|1\rangle$ state, and only the $m_F < 0$ states to the $|0\rangle$ state. This requirement is essential to ensuring that the CNOT does not generate any X- or Y-type errors. For magnetic field regimes typically used in alkali atom Rydberg experiments, this state selectivity may not be straightforward to implement, as the level separation between different m_F states within a single hyperfine manifold may be much smaller than the linewidth of the lasers used for optical pumping. To address this challenge, we can utilize a Rydberg state as a shelving state (due to its long lifetime) to avoid unwanted pumping of $m_F < 0$ (respectively, $m_F > 0$) states to $|1\rangle$ ($|0\rangle$). Thus, in Step 6 of Figure 6.5, the optical pumping of $m_F > 0$ states into the $|1\rangle$ state can be implemented for ^{85}Rb as follows:

1. Swap the population between the $|1\rangle$ state and the stretched ground state $|F = I+1/2, m_F = I + 1/2\rangle$.

2. Swap the population between the $|0\rangle$ state and the ground state $|F = 3, m_F = 0\rangle$.

3. Apply a resonant π pulse to shelve any population in the $|F = 2, m_F = -2\rangle$ state into the Rydberg state $|nS_{1/2}, m_J = -1/2, m_I = -5/2\rangle$.

4. Use σ^+ light to excite states in the $F = 2$ ground state manifold to the $5P_{3/2}$ $F = 3$ manifold; these excited states decay quickly back to the ground state.

5. Apply resonant π pulses $|F = 3, m_F = 1\rangle \leftrightarrow |F = 2, m_F = 1\rangle$ and $|F = 3, m_F = 2\rangle \leftrightarrow |F = 2, m_F = 2\rangle$.

6. Repeat Steps 4 and 5 as necessary; after several iterations, all population that started with $m_F > 0$ will be in the $|F = 3, m_F = 3\rangle$ state.

7. Repeat Steps 1, 2, and 3.

Because the $|F = 2, m_F = -2\rangle$ state can only be populated if a Rydberg error occurred in one of the earlier steps of the bias-preserving CNOT, to leading order in the total error probability, we may assume that the Rydberg state $|nS_{1/2}, m_J = -1/2, m_I = -5/2\rangle$ will not decay if it is populated in the above procedure. In this way, the only $F = 2$ states that can be populated at the beginning of Step 4 above will be the $m_F > 0$ states, so the optical pumping will work in the same way as the protocol described in Section 6.4.2 (Figure 6.6). An analogous procedure can then be applied to pump the $m_F < 0$ states into $|0\rangle$. In this latter case, it will not be necessary to shelve population in the Rydberg state, as all $m_F > 0$ population will already have been transferred to the $|1\rangle$ state.

Bibliography

[1] Aguado, M. & Vidal, G. (2008). Entanglement renormalization and topological order. *Phys. Rev. Lett.*, 100, 070404.

[2] Aharonov, D. & Ben-Or, M. (1997). Fault-tolerant quantum computation with constant error. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing* (pp. 176–188).

[3] Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G., Buell, D. A., et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510.

[4] Auger, J. M., Bergamini, S., & Browne, D. E. (2017). Blueprint for fault-tolerant quantum computation with Rydberg atoms. *Phys. Rev. A*, 96, 052320.

[5] Awschalom, D. D., Bassett, L. C., Dzurak, A. S., Hu, E. L., & Petta, J. R. (2013). Quantum spintronics: Engineering and manipulating atom-like spins in semiconductors. *Science*, 339, 1174–1179.

[6] Bardeen, J., Cooper, L. N., & Schrieffer, J. R. (1957). Theory of superconductivity. *Phys. Rev.*, 108, 1175–1204.

[7] Barredo, D., de Léséleuc, S., Lienhard, V., Lahaye, T., & Browaeys, A. (2016). An atom-by-atom assembler of defect-free two-dimensional atomic arrays. *Science*, 354(6315), 1021–1023.

[8] Barredo, D., Lienhard, V., de Léséleuc, S., Lahaye, T., & Browaeys, A. (2018). Synthetic three-dimensional atomic structures assembled atom by atom. *Nature*, 561, 79–82.

[9] Barredo, D., Lienhard, V., Scholl, P., de Léséleuc, S., Boulier, T., Browaeys, A., & Lahaye, T. (2020). Three-dimensional trapping of individual Rydberg atoms in ponderomotive bottle beam traps. *Phys. Rev. Lett.*, 124, 023201.

[10] Beguin, L. (2013). Measurement of the van der waals interaction between two Rydberg atoms.

[11] Bergschneider, A., Klinkhamer, V. M., Becher, J. H., Klemt, R., Zürn, G., Preiss, P. M., & Jochim, S. (2018). Spin-resolved single-atom imaging of ^{6}Li in free space. *Phys. Rev. A*, 97, 063613.

[12] Bernien, H., Schwartz, S., Keesling, A., Levine, H., Omran, A., Pichler, H., Choi, S., Zibrov, A. S., Endres, M., Greiner, M., & Vuletic, V. (2017). Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682), 579–584.

[13] Bertlmann, R. A. & Krammer, P. (2008). Bloch vectors for qudits. *Journal of Physics A: Mathematical and Theoretical*, 41(23), 235303.

[14] Beterov, I. I., Ryabtsev, I. I., Tretyakov, D. B., & Entin, V. M. (2009). Quasiclassical calculations of blackbody-radiation-induced depopulation rates and effective lifetimes of Rydberg *ns*, *np*, and *nd* alkali-metal atoms with $n \leq 80$. *Phys. Rev. A*, 79, 052504.

[15] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549, 195–202.

[16] Bluvstein, D., Levine, H., Semeghini, G., Wang, T. T., Ebadi, S., Kalinowski, M., Keesling, A., Maskara, N., Pichler, H., Greiner, M., et al. (2022). A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604(7906), 451–456.

[17] Bluvstein, D., Omran, A., Levine, H., Keesling, A., Semeghini, G., Ebadi, S., Wang, T. T., Michailidis, A. A., Maskara, N., Ho, W. W., Choi, S., Serbyn, M., Greiner, M., Vuletic, V., & Lukin, M. D. (2021). Controlling quantum many-body dynamics in driven Rydberg atom arrays. *Science*.

[18] Bombin, H. & Martin-Delgado, M. A. (2006). Topological quantum distillation. *Phys. Rev. Lett.*, 97, 180501.

[19] Bombin, H. & Martin-Delgado, M. A. (2007). Topological computation without braiding. *Phys. Rev. Lett.*, 98, 160502.

[20] Boyd, M. M., Zelevinsky, T., Ludlow, A. D., Blatt, S., Zanon-Willette, T., Foreman, S. M., & Ye, J. (2007). Nuclear spin effects in optical lattice clocks. *Phys. Rev. A*, 76, 022510.

[21] Bravyi, S. & Kitaev, A. (2005). Universal quantum computation with ideal clifford gates and noisy ancillas. *Phys. Rev. A*, 71, 022316.

[22] Bravyi, S. B. & Kitaev, A. Y. (1998). *Quantum codes on a lattice with boundary*. preprint, arXiv.

[23] Bricmont, J. & Frölich, J. (1983). An order parameter distinguishing between different phases of lattice gauge theories with matter fields. *Physics Letters B*, 122(1), 73–77.

[24] Bridgeman, J. C., Flammia, S. T., & Poulin, D. (2016). Detecting topological order with ribbon operators. *Physical Review B*, 94(20), 205123.

[25] Broughton, M., Verdon, G., McCourt, T., Martinez, A. J., Yoo, J. H., Isakov, S. V., Massey, P., Halavati, R., Niu, M. Y., Zlokapa, A., Peters, E., Lockwood, O., Skolik, A., Jerbi, S., Dunjko, V., Leib, M., Streif, M., Von Dollen, D., Chen, H., Cao, S., Wiersema, R., Huang, H.-Y., McClean, J. R., Babbush, R., Boixo, S., Bacon, D., Ho, A. K., Neven, H., & Mohseni, M. (2020). Tensorflow quantum: A software framework for quantum machine learning.

[26] Brown, L. D., Can, T. T., & DasGupta, A. (2001). Interval estimation for a binomial proportion. *Statistical Science*, 16, 101.

[27] Bruzewicz, C. D., Chiaverini, J., McConnell, R., & Sage, J. M. (2019). Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2), 021314.

[28] Brydges, T., Elben, A., Jurcevic, P., Vermersch, B., Maier, C., Lanyon, B. P., Zoller, P., Blatt, R., & Roos, C. F. (2019). Probing Rényi entanglement entropy via randomized measurements. *Science*, 364, 260–263.

[29] Campagne-Ibarcq, P., Eickbusch, A., Touzard, S., Zalys-Geller, E., Frattini, N. E., Sivak, V. V., Reinhold, P., Puri, S., Shankar, S., Schoelkopf, R. J., et al. (2020). Quantum error correction of a qubit encoded in grid states of an oscillator. *Nature*, 584(7821), 368–372.

[30] Carleo, G. & Troyer, M. (2017). Solving the quantum many-body problem with artificial neural networks. *Science*, 355, 602–606.

[31] Carmichael, H. J. (1993). *An Open Systems Approach to Quantum Optics*. Springer.

[32] Carrasquilla, J. & Melko, R. G. (2017). Machine learning phases of matter. *Nat. Phys.*, 13, 431–434.

[33] Carruthers, P. & Nieto, M. M. (1965). Coherent states and the forced quantum oscillator. *American Journal of Physics*, 33(7), 537–544.

[34] Castelnovo, C. & Chamon, C. (2008). Quantum topological phase transition at the microscopic level. *Physical Review B*, 77(5), 054433.

[35] Chamberland, C. & Beverland, M. E. (2018). Flag fault-tolerant error correction with arbitrary distance codes. *Quantum*, 2, 53.

[36] Chamberland, C., Noh, K., Arrangoiz-Arriola, P., Campbell, E. T., Hann, C. T., Iverson, J., Putterman, H., Bohdanowicz, T. C., Flammia, S. T., Keller, A., Refael, G., Preskill, J., Jiang, L., Safavi-Naeini, A. H., Painter, O., & BrandÃ£o, F. G. S. L. (2020). *Building a fault-tolerant quantum computer using concatenated cat codes*. preprint.

[37] Chao, R. & Reichardt, B. W. (2018a). Fault-tolerant quantum computation with few qubits. *npj Quantum Information*, 4, 42.

[38] Chao, R. & Reichardt, B. W. (2018b). Quantum error correction with only two extra qubits. *Phys. Rev. Lett.*, 121, 050502.

[39] Chao, R. & Reichardt, B. W. (2020). Flag fault-tolerant error correction for any stabilizer code. *PRX Quantum*, 1(1), 010302.

[40] Chen, X., Gu, Z.-C., & Wen, X.-G. (2010). Local unitary transformation, long-range quantum entanglement, wave function renormalization, and topological order. *Physical Review B*, 82(15), 155138.

[41] Chen, X., Gu, Z.-C., & Wen, X.-G. (2011). Classification of gapped symmetric phases in one-dimensional spin systems. *Phys. Rev. B*, 83, 035107.

[42] Cheng, Y., Li, C., & Zhai, H. (2021). Variational approach to quantum spin liquid in a Rydberg atom simulator.

[43] Cirac, J. I., Pérez-García, D., Schuch, N., & Verstraete, F. (2021). Matrix product states and projected entangled pair states: Concepts, symmetries, theorems. *Rev. Mod. Phys.*, 93, 045003.

[44] Clarke, J. & Wilhelm, F. K. (2008). Superconducting quantum bits. *Nature*, 453(7198), 1031–1042.

[45] Cohen, J. S. (2003). Effect of tunneling on ionization of Rydberg states in intense fields: Hydrogenic atoms. *Phys. Rev. A*, 68, 033409.

[46] Cooke, W. E. & Gallagher, T. F. (1980). Effects of blackbody radiation on highly excited atoms. *Phys. Rev. A*, 21, 588.

[47] de Neeve, B., Nguyen, T. L., Behrle, T., & Home, J. (2020). *Error correction of a logical grid state qubit by dissipative pumping*. preprint, arXiv.

[48] Dennis, E., Kitaev, A., Landahl, A., & Preskill, J. (2002). Topological quantum memory. *Journal of Mathematical Physics*, 43(9), 4452–4505. Publisher: American Institute of Physics.

[49] Devoret, M. H. & Schoelkopf, R. J. (2013). Superconducting circuits for quantum information: An outlook. *Science*, 339, 1169–1174.

[50] Duclos-Cianci, G. & Poulin, D. (2013). Fault-tolerant renormalization group decoder for abelian topological codes.

[51] Duivenvoorden, K., Iqbal, M., Haegeman, J., Verstraete, F., & Schuch, N. (2017). Entanglement phases as holographic duals of anyon condensates. *Phys. Rev. B*, 95, 235119.

[52] Dunjko, V., Taylor, J. M., & Briegel, H. J. (2016). Quantum-enhanced machine learning. *Phys. Rev. Lett.*, 117, 130501.

[53] Ebadi, S., Wang, T. T., Levine, H., Keesling, A., Semeghini, G., Omran, A., Bluvstein, D., Samajdar, R., Pichler, H., Ho, W. W., Choi, S., Sachdev, S., Grener, M., Vuletic, V., & Lukin, M. D. (2020). *Quantum Phases of Matter on a 256-Atom Programmable Quantum Simulator*. preprint, arXiv.

[54] Egan, L., Debroy, D. M., Noel, C., Risinger, A., Zhu, D., Biswas, D., Newman, M., Li, M., Brown, K. R., Cetina, M., & Monroe, C. (2021). *Fault-Tolerant Operation of a Quantum Error-Correction Code*. preprint, arXiv.

[55] Farhi, E. & Neven, H. (2018). *Classification with Quantum Neural Networks on Near Term Processors*. preprint, arXiv.

[56] Ferain, I., Colinge, C. A., & Colinge, J.-P. (2011). Multigate transistors as the future of classical metal–oxide–semiconductor field-effect transistors. *Nature*, 479(7373), 310–316.

[57] Fowler, A. G., Devitt, S. J., & Hollenberg, L. C. L. (2004). Implementation of shor's algorithm on a linear nearest neighbour qubit array. *Quantum Info. Comput.*, 4(4), 237–251.

[58] Fowler, A. G., Mariantoni, M., Martinis, J. M., & Cleland, A. N. (2012). Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86, 032324.

[59] Fredenhagen, K. & Marcu, M. (1983). Charged states in $\mathbb{Z}_2$ gauge theories. *Communications in Mathematical Physics*, 92(1), 81–119.

[60] Freeman, R. (1998). Shaped radiofrequency pulses in high resolution NMR. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 32(1), 59–106.

[61] Ge, Y., Molnar, A., & Cirac, J. I. (2016). Rapid adiabatic preparation of injective projected entangled pair states and gibbs states. *Phys. Rev. Lett.*, 116, 080503.

[62] Ginzburg, V. L. & Landau, L. D. (2009). On the theory of superconductivity. In *On superconductivity and superfluidity* (pp. 113–137). Springer.

[63] Giudici, G., Lukin, M. D., & Pichler, H. (2022). Dynamical preparation of quantum spin liquids in Rydberg atom arrays.

[64] Gottesman, D. (2010). An introduction to quantum error correction and fault-tolerant quantum computation. In *Quantum Information Science and its Contributions to Mathematics, Proc. Symposium in Applied Math.* (pp. 13–56). Providence, RI: American Mathematical Society.

[65] Graham, T. M., Kwon, M., Grinkemeyer, B., Marra, Z., Jiang, X., Lichtman, M. T., Sun, Y., Ebert, M., & Saffman, M. (2019). Rydberg-mediated entanglement in a two-dimensional neutral atom qubit array. *Phys. Rev. Lett.*, 123, 230501.

[66] Gregor, K., Huse, D. A., Moessner, R., & Sondhi, S. L. (2011). Diagnosing deconfinement and topological order. *New Journal of Physics*, 13(2), 025009.

[67] Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (pp. 212–219).

[68] Guillaud, J. & Mirrahimi, M. (2019). Repetition cat qubits for fault-tolerant quantum computation. *Phys. Rev. X*, 9, 041053.

[69] Haah, J. (2016). An invariant of topologically ordered states under local unitary transformations. *Communications in Mathematical Physics*, 342(3), 771–801. Publisher: Springer.

[70] Haah, J., Harrow, A. W., Ji, Z., Wu, X., & Yu, N. (2017). Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, 63(9), 5628–5641.

[71] Haegeman, J., Pérez-García, D., Cirac, I., & Schuch, N. (2012). Order parameter for symmetry-protected topological phases in one dimension. *Phys. Rev. Lett.*, 109, 050402.

[72] Häffner, H., Roos, C. F., & Blatt, R. (2008). Quantum computing with trapped ions. *Physics reports*, 469(4), 155–203.

[73] Haldane, F. D. M. (1983). Nonlinear field theory of large-spin heisenberg antiferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis néel state. *Phys. Rev. Lett.*, 50, 1153.

[74] Harris, R. et al. (2018). Phase transitions in a programmable quantum spin glass simulator. *Science*, 361(6398), 162–165.

[75] Hastings, M. B. & Wen, X.-G. (2005). Quasiadiabatic continuation of quantum states: The stability of topological ground-state degeneracy and emergent gauge invariance. *Phys. Rev. B*, 72(4), 045141. Publisher: American Physical Society.

[76] Hayes, D., Stack, D., Bjork, B., Potter, A. C., Baldwin, C. H., & Stutz, R. P. (2020). Eliminating leakage errors in hyperfine qubits. *Phys. Rev. Lett.*, 124, 170501.

[77] Herrmann, J., Llima, S. M., Remm, A., Zapletal, P., McMahon, N. A., Scarato, C., Swiadek, F., Andersen, C. K., Hellings, C., Krinner, S., Lacroix, N., Lazar, S., Kerschbaum, M., Zanuz, D. C., Norris, G. J., Hartmann, M. J., Wallraff, A., & Eichler, C. (2022). Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases. *Nature Communications*, 13(1).

[78] Hinton, G. (2007). Lecture Notes for CSC2515: Introduction to Machine Learning, University of Toronto.

[79] Horodecki, M., Horodecki, P., & Horodecki, R. (1996). Separability of mixed states: necessary and sufficient conditions. *Physics Letters A*, 223(1), 1–8.

[80] Huang, C.-Y., Chen, X., & Lin, F.-L. (2013). Symmetry-protected quantum state renormalization. *Phys. Rev. B*, 88, 205124.

[81] Huggins, W., Patil, P., Mitchell, B., Whaley, K. B., & Stoudenmire, E. M. (2018). Towards quantum machine learning with tensor networks. *Quantum Science and Technology*, 4, 024001.

[82] Iqbal, M. & Schuch, N. (2021). Entanglement Order Parameters and Critical Behavior for Topological Phase Transitions and Beyond. *Physical Review X*, 11(4), 041014.

[83] Isenhower, L., Saffman, M., & Mølmer, K. (2011). Multibit C_kNOT quantum gates via Rydberg blockade. *Quantum Inf. Process.*, 10, 755–770.

[84] Jaksch, D., Cirac, J. I., Zoller, P., Rolston, S. L., Cote, R., & Lukin, M. D. (2000). Fast quantum gates for neutral atoms. *Phys. Rev. Lett.*, 85, 2208.

[85] Jamadagni, A., Kazemi, J., & Weimer, H. (2022). Learning of error statistics for the detection of quantum phases.

[86] Kamta, G. L., Engo, S. G. N., Njock, M. G. K., & Oumarou, B. (1998). Consistent description of Klein-Gordon dipole matrix elements. *J. Phys. B.: AT. Mol. Opt. Phys.*, 31, 963.

[87] Kaulakys, B. (1995). Consistent analytical approach for the quasi-classical radial dipole matrix elements. *J. Phys. B: At. Mol. Opt. Phys.*, 28, 4963.

[88] Keesling, A., Omran, A., Levine, H., Bernien, H., Pichler, H., Choi, S., Samajdar, R., Schwartz, S., Silvi, P., Sachdev, S., Zoller, P., Endres, M., Greiner, M., Vuletic, V., & Lukin, M. D. (2019). Quantum Kibble-Zurek mechanism and critical dynamics on a programmable Rydberg simulator. *Nature*, 568, 207–211.

[89] Kim, H., Park, Y. J., Kim, K., Sim, H.-S., & Ahn, J. (2018). Detailed balance of thermalization dynamics in Rydberg-atom quantum simulators. *Phys. Rev. Lett.*, 120, 180502.

[90] Kim, I. & Swingle, B. (2017). *Robust entanglement renormalization on a noisy quantum computer*. preprint, arXiv.

[91] Kitaev, A. & Preskill, J. (2006). Topological Entanglement Entropy. *Phys. Rev. Lett.*, 96(11), 110404. Publisher: American Physical Society.

[92] Kitaev, A. Y. (2003). Fault-tolerant quantum computation by anyons. *Ann. Phys.*, 303, 2–30.

[93] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*.

[94] Labuhn, H., Barredo, D., Ravets, S., de D. Léséleuc, S., Macri, M., Lahaye, T., & Browaeys, A. (2016). Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum ising models. *Nature*, 534(7609), 667–70.

[95] Ladd, T. D. et al. (2010). Quantum computers. *Nature*, 464, 45–53.

[96] Landau, L. D. (1937). On the theory of phase transitions. i. *Zh. Eksp. Teor. Fiz.*, 11, 19.

[97] Laughlin, R. B. (1983). Anomalous quantum hall effect: An incompressible quantum fluid with fractionally charged excitations. *Phys. Rev. Lett.*, 50, 1395–1398.

[98] LeCun, Y. & Bengio, Y. (1995). Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361, 10.

[99] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444.

[100] Lee, J. Y. & Landon-Cardinal, O. (2019). Practical variational tomography for critical one-dimensional systems. *Phys. Rev. A*, 91, 062128.

[101] Lee, Y. A. & Vidal, G. (2013). Entanglement negativity and topological order. *Physical Review A*, 88(4), 042318.

[102] Levin, M. & Wen, X.-G. (2006). Detecting Topological Order in a Ground State Wave Function. *Phys. Rev. Lett.*, 96(11), 110405. Publisher: American Physical Society.

[103] Levin, M. A. & Wen, X.-G. (2005). String-net condensation: A physical mechanism for topological phases. *Phys. Rev B.*, 71, 045110.

[104] Levine, H., Keesling, A., Omran, A., Bernien, H., Schwartz, S., Zibrov, A. S., Endres, M., Greiner, M., Vuletic, V., & Lukin, M. D. (2018). High-fidelity control and entanglement of Rydberg atom qubits. *Phys. Rev. Lett.*, 121, 123603.

[105] Levine, Y., Cohen, N., & Shashua, A. (2019). Quantum entanglement in deep learning architectures. *Phys. Rev. Lett.*, 122, 065301.

[106] Lin, H. W., Tegmark, M., & Rolnick, D. (2017). Why does deep and cheap learning work so well? *J. Stat. Phys.*, 168(6), 1223–1247.

[107] Löw, R., Weimer, H., Nipper, J., Balewski, J. B., Butscher, B., Büchler, H. P., & Pfau, T. (2012). An experimental and theoretical guide to strongly interacting Rydberg gases. *J. Phys. B.: At. Mol. Opt. Phys.*, 45, 113001.

[108] Ludlow, A. D., Boyd, M. M., Ye, J., Peik, E., & Schmidt, P. O. (2015). Optical atomic clocks. *Rev. Mod. Phys.*, 87, 637–701.

[109] Lukin, M. D., Fleischhauer, M., Cote, R., Duan, L. M., Jaksch, D., Cirac, J. I., & Zoller, P. (2001). Dipole blockade and quantum information processing in mesoscopic atomic ensembles. *Phys. Rev. Lett.*, 87, 037901.

[110] Madjarov, I. S., Covey, J. P., Shaw, A. L., Choi, J., Kale, A., Cooper, A., Picher, H., Schkolnik, V., Williams, J. R., & Endres, M. (2020). High-fidelity entanglement and detection of alkaline-earth Rydberg atoms. *Nat. Phys.*

[111] Maller, K. M., Lichtman, M. T., Xia, T., Sun, Y., Piotrowicz, M. J., Carr, A. W., Isenhower, L., & Saffman, M. (2015). Rydberg-blockade controlled-not gate and entanglement in a two-dimensional array of neutral-atom qubits. *Phys. Rev. A*, 92, 022336.

[112] Maskara, N., Kubica, A., & Jochym-O'Connor, T. (2019). Advantages of versatile neural-network decoding for topological codes. *Phys. Rev. A*, 99, 052351.

[113] McCarthy, B. & Ponedal, S. (2021). Ibm unveils world's first 2 nanometer chip technology, opening a new frontier for semiconductors. *IBM News Room.*, (pp. 6–8).

[114] McCulloch, I. P. (2008). Infinite size density matrix renormalization group, revisited. *arXiv preprint arXiv:0804.2509*.

[115] Mehta, P. & Schwab, D. J. (2014). *An exact mapping between the variational renormalization group and deep learning*. preprint, arXiv.

[116] Michael, M. H., Silveri, M., Brierley, R. T., Albert, V. V., Salmilehto, J., Jiang, L., & Girvin, S. M. (2016). New class of quantum error-correcting codes for a bosonic mode. *Phys. Rev. X*, 6, 031006.

[117] Misguich, G., Serban, D., & Pasquier, V. (2002). Quantum Dimer Model on the Kagome Lattice: Solvable Dimer-Liquid and Ising Gauge Theory. *Phys. Rev. Lett.*, 89(13), 137202. Publisher: American Physical Society.

[118] Monroe, C. & Kim, J. (2013). Scaling the ion trap quantum processor. *Science*, 339, 1164–1169.

[119] Moore, G. E. (1965). Cramming more components onto integrated circuits.

[120] Napp, J. C., La Placa, R. L., Dalzell, A. M., Brandão, F. G. S. L., & Harrow, A. W. (2022). Efficient classical simulation of random shallow 2d quantum circuits. *Phys. Rev. X*, 12, 021021.

[121] Nayak, C., Simon, S. H., Stern, A., Freedman, M., & Das Sarma, S. (2008). Non-abelian anyons and topological quantum computation. *Rev. Mod. Phys.*, 80, 1083–1159.

[122] Nielsen, M. A. & Chuang, I. (2000). *Quantum Computation and Quantum Information*. Cambridge University Press.

[123] Nigg, D., Müller, M., Martinez, E. A., Schindler, P., Hennrich, M., Monz, T., Martin-Delgado, M. A., & Blatt, R. (2014). Quantum computations on a topologically encoded qubit. *Science*, 345(6194), 302–305.

[124] Omran, A., Levine, H., Keesling, A., Semeghini, G., Wang, T. T., Ebadi, S., Bernien, H., Zibrov, A. S., Pichler, H., Choi, S., Cui, J., Rossignolo, M., Rembold, P., Montangero, S., Calarco, T., Endres, M., Greiner, M., Vuletic, V., & Lukin, M. D. (2019). Generation and manipulation of Schrödinger cat states in Rydberg atom arrays. *Science*, 365(6543), 570–574.

[125] Peres, A. (1996). Separability Criterion for Density Matrices. *Physical Review Letters*, 77(8), 1413–1415.

[126] Pfeifer, R. N. C., Evenbly, G., & Vidal, G. (2009). Entanglement renormalization, scale invariance, and quantum criticality. *Phys. Rev. A*, 79, 040301.

[127] Poilblanc, D., Schuch, N., Pérez-García, D., & Cirac, J. I. (2012). Topological and entanglement properties of resonating valence bond wave functions. *Phys. Rev. B*, 86(1), 014404. Publisher: American Physical Society.

[128] Pollmann, F. & Turner, A. M. (2012). Detection of symmetry-protected topological phases in one dimension. *Phys. Rev. B*, 86, 125441.

[129] Preskill, J. (1998). Quantum information and computation. Lecture Notes for Physics 229, California Institute of Technology.

[130] Puri, S., St-Jean, L., Gross, J. A., Grimm, A., Frattini, N. E., Iyer, P. S., Krishna, A., Touzard, S., Jiang, L., Blais, A., Flammia, S. T., & Girvin, S. M. (2019). *Bias-preserving gates with stabilized cat qubits*. preprint, arXiv.

[131] Reichardt, B. W. (2018). *Fault-tolerant quantum error correction for Steane's seven-qubit color code with few or no extra qubits*. preprint, arXiv.

[132] Rivest, R. L., Shamir, A., & Adleman, L. (1978). A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2), 120–126.

[133] Rosenblum, S., Reinhold, P., Mirrahimi, M., Jiang, L., Frunzio, L., & Schoelkopf, R. J. (2018). Fault-tolerant detection of a quantum error. *Science*, 361(6399), 266–270.

[134] Sachdev, S. (1992). Kagome- and triangular-lattice Heisenberg antiferromagnets: Ordering from quantum fluctuations and quantum-disordered ground states with unconfined bosonic spinons. *Phys. Rev. B*, 45(21), 12377–12396. Publisher: American Physical Society.

[135] Sachdev, S. (2011). *Quantum Phase Transitions*. Cambridge University Press.

[136] Saffman, M., Walker, T. G., & Mølmer, K. (2010). Quantum information with Rydberg atoms. *Rev. Mod. Phys.*, 82, 2313.

[137] Samajdar, R., Joshi, D. G., Teng, Y., & Sachdev, S. (2022). Emergent $\mathbb{Z}_2$ gauge theories and topological excitations in Rydberg atom arrays.

[138] Satzinger, K., Liu, Y.-J., Smith, A., Knapp, C., Newman, M., Jones, C., Chen, Z., Quintana, C., Mi, X., Dunsworth, A., et al. (2021). Realizing topologically ordered states on a quantum processor. *Science*, 374(6572), 1237–1241.

[139] Schuch, N., Pérez-García, D., & Cirac, I. (2011). Classifying quantum phases using matrix product states and projected entangled pair states. *Phys. Rev. B*, 84, 165139.

[140] Schuch, N., Poilblanc, D., Cirac, J. I., & Pérez-García, D. (2012). Resonating valence bond states in the PEPS formalism. *Phys. Rev. B*, 86(11), 115108. Publisher: American Physical Society.

[141] Schuch, N., Wolf, M. M., Verstraete, F., & Cirac, J. I. (2007). Computational complexity of projected entangled pair states. *Phys. Rev. Lett.*, 98, 140506.

[142] Schwarz, M., Temme, K., & Verstraete, F. (2012). Preparing projected entangled pair states on a quantum computer. *Phys. Rev. Lett.*, 108, 110502.

[143] Semeghini, G., Levine, H., Keesling, A., Ebadi, S., Wang, T. T., Bluvstein, D., Verresen, R., Pichler, H., Kalinowski, M., Samajdar, R., et al. (2021). Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572), 1242–1247.

[144] Shi, Y. (2003). Both Toffoli and Controlled-NOT need little help to do universal quantum computation. *Quantum Information and Computation*, 3, 84–92.

[145] Shor, P. W. (1994). Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th annual symposium on foundations of computer science* (pp. 124–134).: Ieee.

[146] Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A*, 52, R2493–R2496.

[147] Shor, P. W. (1996). Fault-tolerant quantum computation. In *Proceedings of 37th Conference on Foundations of Computer Science* (pp. 56–65).: IEEE.

[148] Singh, K., Anand, S., Pocklington, A., Kemp, J. T., & Bernien, H. (2021). *A dual-element, two-dimensional atom array with continuous-mode operation*. preprint, arXiv.

[149] Singh, S. & Vidal, G. (2013). Symmetry-protected entanglement renormalization. *Phys. Rev. B*, 88, 121108(R).

[150] Steane, A. (1996). Multiple-particle interference and quantum error correction. *Proc. R. Soc. Lond. A.*, 452, 2551–2577.

[151] Steck, D. A. (2019). Rubidium 87 D line data. available online at http://steck.us/alkalidata (revision 2.2.1, 21 November 2019).

[152] Stormer, H. L., Tsui, D. C., & Gossard, A. C. (1999). The fractional quantum hall effect. *Rev. Mod. Phys.*, 71, S298–S305.

[153] Stricker, R., Vodola, D., Erhard, A., Postler, L., Meth, M., Ringbauer, M., Schindler, P., Monz, T., Müller, M., & Blatt, R. (2020). Experimental deterministic correction of qubit loss. *Nature*, 585(7824), 207–210.

[154] Tarabunga, P. S., Surace, F. M., Andreoni, R., Angelone, A., & Dalmonte, M. (2022). Gauge-theoretic origin of rydberg quantum spin liquids.

[155] Terhal, B. M. (2015). Quantum error correction for quantum memories. *Rev. Mod. Phys.*, 87, 307–346.

[156] Tuckett, D. K., Bartlett, S. D., & Flammia, S. T. (2018). Ultrahigh error threshold for surface codes with biased noise. *Phys. Rev. Lett.*, 120, 050505.

[157] van Nieuwenburg, E. P. L., Liu, Y. H., & Huber, S. D. (2017). Learning phase transitions by confusion. *Nat. Phys.*, 13, 435–439.

[158] Vandersypen, L. M. K. & Chuang, I. L. (2005). Nmr techniques for quantum control and computation. *Rev. Mod. Phys.*, 76, 1037–1069.

[159] Vedral, V., Barenco, A., & Ekert, A. (1996). Quantum networks for elementary arithmetic operations. *Phys. Rev. A*, 54, 147–153.

[160] Verresen, R., Lukin, M. D., & Vishwanath, A. (2021a). Prediction of toric code topological order from Rydberg blockade. *Phys. Rev. X*, 11, 031005.

[161] Verresen, R., Moessner, R., & Pollmann, F. (2017). One-dimensional symmetry protected topological phases and their transitions. *Phys. Rev. B*, 96, 165124.

[162] Verresen, R., Tantivasadakarn, N., & Vishwanath, A. (2021b). Efficiently preparing SchrÃ¶dinger's cat, fractons and non-abelian topological order in quantum devices.

[163] Verresen, R. & Vishwanath, A. (2022). Unifying Kitaev magnets, kagome dimer models and ruby Rydberg spin liquids.

[164] Vidal, G. (2003). Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.*, 91, 147902.

[165] Vidal, G. (2008a). Class of many-body states that can be efficiently simulated. *Phys. Rev. Lett.*, 101, 110501.

[166] Vidal, G. (2008b). Class of quantum many-body states that can be efficiently simulated. *Phys. Rev. Lett.*, 101, 110501.

[167] Wang, L. (2016). Discovering phase transitions with unsupervised learning. *Phys. Rev. B*, 94, 195105.

[168] Wang, P., Luan, C.-Y., Qiao, M., Um, M., Zhang, J., Wang, Y., Yuan, X., Gu, M., Zhang, J., & Kim, K. (2021). Single ion qubit with estimated coherence time exceeding one hour. *Nature Communications*, 12(1), 1–8.

[169] Wen, X.-G. (2017). Colloquium: Zoo of quantum-topological phases of matter. *Rev. Mod. Phys.*, 89(4), 041004. Publisher: American Physical Society.

[170] Wilson, J., Saskin, S., Meng, Y., Ma, S., Dilip, R., Burgers, A., & Thompson, J. (2019). *Trapped arrays of alkaline earth Rydberg atoms in optical tweezers*. preprint, arXiv.

[171] Wilson, K. G. (1974). Confinement of quarks. *Physical Review D*, 10(8), 2445–2459.

[172] Wootton, J. R. (2012). A witness for topological order and stable quantum memories in abelian anyonic systems. *Journal of Physics A: Mathematical and Theoretical*, 45(39), 395301.

[173] Wu, Y., Kolkowitz, S., Puri, S., & Thompson, J. D. (2022). *Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays*. preprint, arXiv.

[174] Xu, W., Venkatramani, A. V., Cantu, S. H., Šumarac, T., Klüsener, V., Lukin, M. D., & Vuletić, V. (2021). *Fast Preparation and Detection of a Rydberg Qubit using Atomic Ensembles*. preprint, arXiv.

[175] Yoder, T. J. & Kim, I. H. (2017). The surface code with a twist. *Quantum*, 1, 2.

[176] Yoder, T. J., Takagi, R., & Chuang, I. L. (2016). Universal fault-tolerant gates on concatenated stabilizer codes. *Phys. Rev. X*, 6, 031039.

[177] Zalka, C. (1998). Fast versions of Shor's quantum factoring algorithm. *arXiv preprint quant-ph/9806084*.

[178] Zeng, B. & Zhou, D. L. (2016). Topological and error-correcting properties for symmetry-protected topological order. *Eur. Phys. Lett.*, 113, 56001.

[179] Zeng, Y., Xu, P., He, X., Liu, Y., Liu, M., Wang, J., Papoular, D. J., Shlyapnikov, G. V., & Zhan, M. (2017). Entangling two individual atoms of different isotopes via Rydberg blockade. *Phys. Rev. Lett.*, 119, 160502.

[180] Zhang, J., Pagano, G., Hess, P. W., Kyprianidis, A., Becker, P., Kaplan, H., Gorshkov, A. V., Gong, Z.-X., & Monroe, C. (2017). Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature*, 551(7682), 601–604.

[181] Zhang, S., Lu, Y., Zhang, K., Chen, W., Li, Y., Zhang, J.-N., & Kim, K. (2020). Error-mitigated quantum gates exceeding physical fidelities in a trapped-ion system. *Nature communications*, 11(1), 1–8.

[182] Zhang, Y. & Kim, E.-A. (2017). Quantum loop topography for machine learning. *Phys. Rev. Lett.*, 118, 216401.

[183] Zhong, H.-S., Wang, H., Deng, Y.-H., Chen, M.-C., Peng, L.-C., Luo, Y.-H., Qin, J., Wu, D., Ding, X., Hu, Y., et al. (2020). Quantum computational advantage using photons. *Science*, 370(6523), 1460–1463.