Thermal Effects in Topological Systems

Christopher N. Self

Imperial College London

Department of Physics

July 19, 2018

Supervisors: Dr. Jiannis K. Pachos & Prof. Terry Rudolph

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy from Department of Physics of Imperial College London

Declarations

The work in this thesis was carried out by the author along with the following main collaborators. Chapter 3 was carried out in collaboration with James R. Wootton and Benjamin J. Brown. The work presented in chapter 5 was a collaboration with Sofyan Iblisdir and James R. Wootton. Chapter 6 is a collaboration with Sofyan Iblisdir, with helpful suggestions made by Chris J. Turner. The ongoing work documented in appendix C is being carried out in collaboration with Johannes Knolle and Sofyan Iblisdir. All work was carried out in collaboration with and under the guidance of my supervisor Jiannis K. Pachos. The computational work was carried out by me alone. I certify that the presented material is my own except where otherwise acknowledged and that the work of others is appropriately referenced.

The numerical results of chapter 3 were obtained from simulations run on the High Performance Computing Facilities of the Imperial College Research Computing Services, DOI: 10.14469/hpc/2232. The numerical work included in chapters 5 and 6 was undertaken on ARC2, part of the High Performance Computing Facilities at the University of Leeds, UK.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work

To my parents

Acknowledgements

I am lucky to have worked and mixed with so many excellent scientists. Throughout my PhD my supervisor Jiannis Pachos has helped guide me to interesting problems and provided insights about how to solve them and I am grateful to have worked with him. I am also particularly grateful to the fantastic collaborators I have worked closely with: Benjamin J. Brown, James R. Wootton, Sofyan Iblisdir and Johannes Knolle. Beyond them, a number of other people have been instrumental in the development of my understanding of the physics I have studied in particular: Konstantinos Meichanetzidis, Giandomenico Palumbo, James de Lisle, Christopher J. Turner, Alexander Bullivant, Zlatko Papic, Adrian Hutter and Ville Lahtinen. I am indebted to the doctoral training centre on Controlled Quantum Dynamics, of which I am a student, for the excellent environment and many opportunities it has provided me. In particular, I am extremely grateful to the administrative staff in Whiteley 115, Miranda Smith and Lisa Cheung, for their support and guidance (both practical and social). I am also very grateful to Terry Rudolph, the director of the doctoral training centre who co-supervised me, for his advice and mentoring. My time as a student would have also been much less enjoyable without my fellow students of cohort 4: Izzie Rabey, Antony Milne, Kamil Korzekwa, Matteo Lostaglio, Claudio Polisseni, David Murgia, Max Frenzel, Chris Sparrow and Zachary Bluden-Codd.

Student life is not all work and I am very grateful to many dear friends, both inside and outside physics, for all the time we have spent together. Most importantly, I thank my parents who continue to be unreasonably supportive of me, as well as being generous with both their love and the contents of their fridge.

My time as a PhD student has been challenging and rewarding. Finally I would like to acknowledge the work of the Imperial College London and Leeds University student counselling services, who have supported me at different times while I have been a student. Their role is vital and I am very grateful to them.

Abstract

Topological systems are an exciting frontier of condensed matter physics, as well as having deep relationships to quantum computing. Progress continues to be made both towards building practical quantum computers and experimentally identifying topological features in real materials. However, these attempts have to contend with the messiness of the real world including the effects of finite temperatures.

In quantum computing thermal errors are studied through the framework of quantum error correction. This allows a characterisation of how damaging these errors are, and has provided many general results on the limits of how well given quantum codes can endure thermal errors. In the first part of this thesis we study the typical response of one of the best known quantum codes (the two-dimensional toric code) when subjected to thermal noise. This typical behaviour is compared to the limiting bounds and we note differences that arise due to entropic effects.

Topological features are typically properties of the ground state of a system. In some cases this means finite temperatures can be damaging, washing out the topological properties with excitations above the ground state. But in other cases it has been shown that topological systems can have their own unique finite temperature behaviours, which can be used to identify them. In the second part of the thesis we study the topological model Kitaev's honeycomb. Fixing certain conserved quantities restricts the model to special sectors, some of which carry topological phases. We begin by investigating an intrinsically finite temperature property of these topological phases, thermal edge currents. At too high temperatures, close to the bulk gap, the topological behaviour of these currents breaks down. However, we show that at lower temperatures they can be identified and shown to have the expected properties such as robustness to perturbations. Following this we investigate the thermal states of the full model. Our approach to this allows us to probe the topological order of the model in a way that has not been done before and we identify a number of interesting new features of the finite temperature physics.

Contents

Contents **6**

List of Figures **9**

1 Introduction **11**

2 Background 1: Quantum memories at finite temperature **19**

2.1 Stabiliser codes 19

2.1.1 Construction 19

2.1.2 Detecting and correcting errors 23

2.2 Quantum memory Hamiltonians 25

2.2.1 Topological order and toric code anyons 26

2.3 Modelling finite temperatures 28

2.3.1 Deriving the rate equation 29

2.3.2 Coherence time 33

2.3.3 Simulating thermal physics 34

2.4 Decoding errors 35

2.4.1 Clustering decoder 36

2.4.2 Testing for success 38

2.5 Summary 41

3 Understanding coherence times of the toric code **43**

3.1 Thermal excitations and random walks 44

3.1.1 Small system limit 44

3.1.2 Large system limit 50

3.2 Conclusions 53

4 Background 2: Kitaev's Honeycomb **55**

4.1 Spin Model . . . 55

4.2 Exact solution . . . 57

4.2.1 Majorana fermions . . . 58

4.2.2 Kitaev's fermionisation . . . 58

4.2.3 Vortex sectors and gauge fixing . . . 60

4.2.4 Normal modes . . . 62

4.2.5 Topological phases, the toric code and Ising anyons . . . 64

4.3 Summary . . . 67

5 Topological edge currents **68**

5.1 Ideal temperature range . . . 69

5.2 Energy Currents . . . 71

5.2.1 Defining a current operator . . . 71

5.2.2 Central charges . . . 73

5.2.3 Topological robustness . . . 76

5.2.4 Currents at $T = 0$ . . . 77

5.3 Conclusions . . . 78

6 Kitaev's Honeycomb at Finite Temperature **80**

6.1 Thermal states of the fermions . . . 81

6.1.1 Unprojected fermions . . . 84

6.1.2 Projected fermions . . . 85

6.1.3 Equivalent cases . . . 87

6.2 Aside – Ground state topological degeneracy . . . 89

6.3 Thermal distributions of vortices . . . 92

6.3.1 Monte Carlo methods . . . 93

6.4 Finite temperature properties . . . 96

6.4.1 Thermodynamics . . . 97

6.4.2 Fermionic gap . . . 100

6.4.3 Thermal metal behaviour . . . 102

6.5 Conclusions . . . 105

7 Conclusions **108**

References **121**

Appendices **122**

A Kitaev's Honeycomb, Solved in a Matrix Language **123**

A.1 Skew-symmetric matrices 124

A.2 Normal modes 125

A.3 Correlation functions 128

B Numerical notes and methods **130**

B.1 Kitaev's Honeycomb 130

B.2 Analysis of Monte Carlo data 134

C Dynamics of Kitaev's Honeycomb **139**

C.1 Kubo Formula 140

C.2 Proposal: Quantum evolution after a quench 145

List of Figures

2.1 Building blocks of the toric code and their relationships to topology. 21

2.2 Errors and error syndromes in the toric code 24

2.3 Distance metric used by the clustering decoder we implement 37

2.4 Graphic illustration of the steps of the clustering decoding algorithm 39

2.5 Assessing decoder success by making deductions about Pauli frames 40

2.6 Benchmarking our decoder implementation against i.i.d. noise 41

3.1 Small and large system size limits of toric code thermal errors 45

3.2 Dependence of Π (the probability that an anyon pair do not recombine) on system size and temperature in the *small* system size limit 47

3.3 Pair creation times in the *small* system limit of the toric code memory . . . 49

3.4 Anyon motion times in the *small* system limit of the toric code memory . . 49

3.5 Coherence times of the toric code memory in the *large* system limit 52

3.6 Final anyon pair separation in the *large* system limit 52

3.7 System size independence of the *large* system size coherence time 53

4.1 Kitaev Honeycomb model defined in the spin-picture 56

4.2 Transformation from spins to fermions, used to reach an exact solution . . . 59

4.3 Fermionic representation of the model in the no-vortex sector 60

4.4 Phase diagram of the no-vortex sector of the Kitaev honeycomb 65

4.5 Sketch of the lattice regions used to compute the Chern number 66

5.1 Dispersion relation of the topological edge modes 70

5.2 Non-topological divergence of the edge currents from their ideal behaviour . 74

5.3 Dependence of the edge currents on temperature 74

5.4 Edge currents showing a transition between different topological phases . . 75

5.5 Topological robustness of the edge currents 76
5.6 Finite size scaling analysis of the currents at zero temperature 77

6.1 Generating topologically distinct ground states of the Kitaev Honeycomb, and investigating the finite size splitting of the ground state degeneracy . . 91
6.2 Thermal fermion energy in the ground state sector as temperature goes to zero, showing splitting enforced by the parity projection 92
6.3 Illustration of the Markov step proposal used in our numerics 95
6.4 Thermodynamic properties in the toric code and Ising regimes 98
6.5 Fermionic gap as a function of temperature 100
6.6 Finite size scaling of the fermion gap at zero temperature 101
6.7 Finite size scaling of the fermion gap at infinite temperature 102
6.8 Distribution of fermionic energies near zero energy in the Ising regime at high temperatures 104
6.9 Weak multifractality of the fermionic wavefunctions in the Ising regime at high temperatures 105

1.

Introduction

Topological order and quantum computing are deeply interrelated topics [1], both of which are making the journey from theory to reality. In the last two years some of the largest conventional computing companies (including Intel, IBM and Google) have joined a growing field of start-up companies in unveiling small quantum computing prototypes [2]. Some of these are already accessible to the scientific community through online access; allowing at least one researcher to begin developing quantum computer games to run on these devices and playfully test their limits [3]. Despite the progress made, however, many of the scalability problems faced by these small computers and the technologies they are built on are incredibly hard engineering challenges. In response a large part of the theory community continues to focus on studying alternative approaches and searching for fundamentally different physical architectures.

Topologically ordered systems offer one such alternative [4, 5]. Just as the material physics of semiconductors makes them perfect for building classical computers, the properties of topologically ordered materials would make them natural candidates for constructing quantum chips. Currently, however, there are still many gaps in the theoretical and experimental understanding of topological systems. These gaps need to be filled if we hope to be able to control them with the precision needed to carry out quantum computations.

My thesis reflects the close relationship between quantum computing and topological order by being split between the two topics. This is partly a result of the environments I have been lucky enough to be in during my PhD[1], and partly comes down to the kind of work I enjoy doing. The main focus throughout is studying thermal effects in these different settings, and this is mostly carried out using heuristics and numerical data from simulations. Though the focus is the same, the motivations vary. Finite temperatures are known to be very damaging to quantum computers. The first part (chapters 2 and 3) relates to the effect of these thermal errors in a particular quantum computer design,

[1] Imperial College's quantum information group for the first year and Leeds University's topological quantum computing group after that.

the *toric code*. Whereas the following chapters (chapters 4, 5 and 6) study finite temperature phenomena in topologically ordered systems and topological phases, instead with the motivation that these effects might be useful measurable signatures. Since the models studied and tools used in these two parts are quite different I've split the background material between chapters 2 and 4. Each of these background chapters aim to provide a self contained introduction for the work that follows. The rest of this introduction sets up the problems studied and gives a review of the relevant literature.

Chapters 2 & 3

One of the biggest hurdles to practical quantum computing is combating the effects of noise. Quantum computers are highly sensitive to external noise as they rely on quantum superpositions, which can easily decohere into classical mixed states. The main source of decoherence is interaction with an external environment. However some level of coupling to the outside world is unavoidable. This interaction could be intentional, such as when inputting/outputting data or instructing the computer to perform an operation, or it could just be our inability to perfectly shield the device from temperature or external radiation. The breakthrough invention of quantum error correction codes [6–8], combined with fault-tolerance threshold theorems [9], proved that these damaging effects can be handled and corrected for. Additionally, studying such codes provided a framework for doing so. The study of quantum error correction has been highly successful; many subsequent advances have led to the design of extremely computationally powerful codes and robust error correction schemes [10–14]. The language of codes has also proved to be a useful tool in other areas, being applied as far afield as studying black holes [15].

One important goal of quantum error correction is searching for good *quantum memories*. These specialised codes would allow quantum states to be saved reliably, thus enabling execution of arbitrarily long quantum programs. Theoretical work on quantum memories separates between *active* and *passive* schemes. Active codes rely on the repeated measurement of special error check observables [16]. The outcomes of these measurements are processed with classical algorithms, called *decoders*, to choose a correction to apply. This process is carried out in real time in order to stop small errors becoming large and

problematic. Passive schemes, in contrast, try to find systems whose natural internal processes protect encoded information without the need for constant external correction [17]. For example, by encoding information in the degenerate ground states of a gapped Hamiltonian system. Passive quantum memories whose storage lifetimes increase quickly as the volume of the system is increased[2] are called *self-correcting* [18]. Whether self-correcting quantum memories can be found is an important open problem. Besides how useful self-correcting quantum memories would be, their discovery would also have deep foundational implications. They would demonstrate stable complex quantum order in a macroscopic system at finite temperature, breaking the separation we observe between the quantum micro- and classical macro- worlds.

Unfortunately, many of the established results concerning self-correcting quantum memories are negative results. These come in the form of *no-go theorems* [19, 20] and general statements about the decoherence times of memories interacting with generic thermal environments [21–24]. These negative results show that many of the best studied quantum memory proposals cannot be self-correcting. To avoid these no-go results, more exotic models have been explicitly constructed by finding loopholes in their proofs. These achieve better information storage times, though for the models where it is possible to obtain clear results on storage times they only offer a weaker form of self-correction called *partial* self-correction [25, 26]. Partially self-correcting models exploit an interesting connection to classical glassiness to slow down the thermal growth of errors and protect the encoded information. However, their storage times only increase slowly with system size[3] and above a critical size their lifetimes cease to grow. Another interesting positive result is that models defined in higher dimensions can exhibit self-correction [27, 28]. In particular the *toric code* in four spatial dimensions has been shown to be self-correcting [29]. This insight has led to attempts to construct self-correcting memories by embedding higher dimensional code structures into three-dimensional space, using either fractal lattices [30] or other techniques [31]. However, the complexity of these proposals makes them difficult to definitively analyse. An even harder prospect is analysing quantum memory proposals that are not exactly solvable. These are systems whose code properties are known but we

[2]Self-correction implies a storage times that increases exponentially with system size.
[3]Storage times of partially self-correcting quantum memories increase polynomially with system size.

do not have exact expressions for the states of the models [13, 18]. This gap in knowledge takes rigorous study of their storage times beyond what is possible with the tools currently available. It has been argued that some of these proposals may be self-correcting.

An important theme in quantum memories is the interplay between energy and entropy. Studies of passive quantum memories often make use of a heuristic called the *Arrhenius law*. This estimates the lifetime of a memory based only on the energy costs of bad errors that destroy the information encoded in the memory [17]. The minimum energy needed to introduce such an error is called the *energy barrier* of the memory. The Arrhenius law provides an estimate of the typical time we expect to wait for such errors to occur naturally due to thermal fluctuations. Specifically, if the energy barrier is Δ the Arrhenius law says we expect thermal processes to overcome this barrier and destroy the encoded information in time $\tau \sim e^{\Delta/T}$ where T is the temperature. Many no-go results actually prove bounds on these energy barriers [19, 20], for example by proving that they are constants independent of system size. The Arrhenius law would then imply that the lifetime must also be independent of system size. This is backed up by our knowledge of classical memories, where an energy barrier that scales at least linearly with system size seems to be necessary for self-correction. It was hoped that for quantum memories it might be possible to do better than the Arrhenius law and achieve partial self-correction even without a positively scaling energy barrier by using entropic protection [32]. In this approach the geometry of the system is altered to try to 'hide' the low energy error paths. Maximising entropy would then drive the evolution away from these paths toward higher energy routes and increase the coherence time. However, more recent work seems to rule out this possibility. Arrhenius-like relations rigorously linking the scaling of the energy barrier to an upper bound on memory lifetimes have now been proven in a number of cases [22–24].

The Arrhenius law may function well as an upper bound, but it is dangerous to consider it as a true estimate of the actual lifetime a memory will have. Classical examples exist where a direct application of Arrhenius law misses important entropic effects that radically decrease the lifetime of a memory [33].

Chapter 2 introduces the mathematical and numerical tools used to define and study

quantum memories. This includes the construction of codes, the decoding of errors and the modelling of finite temperatures.

In chapter 3 we investigate the actual expected memory lifetime of a well-known quantum memory model, the two-dimensional toric code, and compare it to the Arrhenius law. This is carried out using arguments from statistical mechanics and numerical simulations. The work in this chapter was included in the publication: [17]. A similar study was published by another group around the same time [34]. Our results agree with theirs in one limiting case (low temperature/small systems), but due to different approaches taken in the analysis we disagree in the other limit (high temperature/large systems).

Chapters 4, 5 & 6

Topological systems are an exciting frontier of condensed matter physics. Conventionally the large-scale (emergent) behaviour of a material, described by its *quasiparticle* excitations, has been categorised as either *fermionic* or *bosonic*. This encodes (amongst other things) the exchange statistics, which describe how the quantum state changes if two quasiparticles are swapped. For bosons there is no change at all. For fermions a simple change occurs, which is undone if the quasiparticles are swapped a second time back to their original positions. The quasiparticles of topologically ordered systems extend these possible behaviours by allowing more general changes of the quantum state [35]. These *anyonic* excitations live in two spatial dimensions; crucially this means looping one particle around another (a double swap) is topologically distinct from doing nothing (no swaps). Anyons are sensitive to the topology of these *braid* loops.

How general the state transformations generated by braids can be depends on the type of anyons. The simplest type, *Abelian* anyons, can exhibit measurable interference between particle trajectories that do or do not braid (this is demonstrated in the Aharanov-Bohm effect [36]). Braids of *non-Abelian* anyons are more powerful and are sufficient to change the result of quantum measurements on the anyons, for example their braids can perform quantum teleportation between subsets of the particles [35]. The most powerful subtypes of non-Abelian anyons are even known to be *universal* for quantum computing [37, 38], meaning their braids alone can carry out any computational operation we would need.

Discovering a physical system with these anyons, that we could precisely control, would provide a strong physical basis for a future quantum computer.

Besides this connection, there is a deep relationship between topological order and quantum error correction. In contrast to conventional phases of matter (such as solids, liquids and gases, or even more exotic phases such as superconductors), topologically ordered systems do not have a *local order parameter*. For a system with solid and gas phases in order to find out what phase the system is in we only need to carry out measurements on a small part, for example by measuring the density. The same is not true for topologically ordered systems. In order to fully characterise their state you need to access the entire system. This confers natural robustness against noise and perturbations, since locally acting errors cannot decohere the non-locally encoded quantum state [17].

Another example of topological states of matter are topological phases [39]. These systems do not directly support anyonic excitations. Nevertheless, they still have important links to topological quantum computing. Though their quasiparticles are not anyons, physical defects in the system (such as vortices) can bind anyons [40]. As with topologically ordered systems, topological phases do not have local order parameters. Instead they are characterised by *topological invariants*. These invariants, such as the *Chern number* in two-dimensional systems [41], are properties of their ground states. Systems with topological phases, such as topological insulators [42] and topological superconductors [43], have gapped energy spectra making them insulators. In spite of this, however, they support robust currents on their surfaces.

There is strong experimental evidence for the existence of topological phases in real materials [42, 43]. Beyond this, it has also been possible to engineer topological phases in quantum simulators using cold atoms [44–46]. In contrast, direct evidence of topological order remains elusive. There are compelling indications that some real setups could carry anyonic excitations, in particular *fractional quantum Hall* systems [47–49] and crystals thought to be *spin liquids* [50]. However, no-one has yet conclusively demonstrated braiding of real quasiparticles [40]. It is hoped that this will eventually be possible, or alternatively that quantum simulator schemes could be extended to produce topologically ordered states.

An important theoretical model of topological matter is Kitaev's Honeycomb [51]. This model is the subject of chapters 4, 5 and 6. It is a two-dimensional spin model that is exactly solvable, meaning it is possible to carry out a deep analysis of its properties. The model has been shown to exhibit both topological order [52] and topological phases [51, 53, 54]. Chapter 4 gives the definitions of the model and describes its exact solution. It ends with a brief dip into a tiny part of the (extensive) literature around the model's properties. This chapter is supported by appendix A, which provides an extended discussion of some of the details of the calculation.

The compelling and beautiful thing about Kitaev's Honeycomb is that its tractability allows us to see the abstract properties (such as its anyons) emerging from the microscopic details. In this vein chapter 5 studies the surface currents (in this case one-dimensional edge currents) of the topological phases of the model using its exact microscopic solution. The topological phases of the Honeycomb are topological superconductors, meaning charge is not conserved and electrical currents are not possible. Instead the currents we study are heat currents, which we study numerically. This is carried out in real-space allowing us to demonstrate their topological robustness. Additionally, the real-space nature of the operators could make it possible to experimentally probe the same observables in future quantum simulator realisations of the model. This work was published as [55].

It is thought that certain crystals may come close to realising the couplings of Kitaev's Honeycomb [56, 57]. Continuing experiments on these materials are able to probe them in ever greater depth [50, 58] leading to a recent increased interest in measurable signatures of Kitaev honeycomb physics. Its thermodynamic properties are one such set of measurable quantities. They have already received a great deal of attention [59–63]. A key finding has been a signature of the *fractionalisation* of the spins in the behaviour of the specific heat. Fractionalisation implies the presence of anyons, but cannot distinguish between different types of anyons. Despite this success however, these previous studies have neglected an important aspect of the topological order of the model that may help to further distinguish the anyon types. Chapter 6 returns to the thermodynamics of the Kitaev Honeycomb, paying close attention to these fine details.

Chapter 7 concludes by summarising the results of chapters 3, 5 and 6. The results

presented are contextualised against other work in the field. After summing up, I end by discussing future directions.

Though the subject matter varies the methods applied throughout the thesis are similar, mostly relying on numerical simulations. I have included an appendix of 'experimental' details on these numerics, appendix B, and much of the original code is available on my github profile @chris-n-self.

Finally, appendix C documents work in progress that seeks to link chapters 5 and 6 by computing thermal conductivities for the full Kitaev Honeycomb model. Two schemes are proposed for computing these quantities. One of these schemes involves computing heat currents after a quenched perturbation, and may separately be of interest in probing aspects of the non-linear response.

2.

Background 1:
Quantum memories at finite temperature

This chapter provides the background on quantum memories and their interaction with a thermal environment. It begins with an introduction to stabiliser codes in section 2.1. Their construction is discussed in section 2.1.1 and the handling of errors in section 2.1.2. The toric code is used as a particular example of a stabiliser code throughout this section and its definitions are given alongside the more general ones. Following this, quantum memory Hamiltonians based on stabiliser codes are introduced in section 2.2. Section 2.2.1 comments briefly on the relationship between quantum memories and topological order. Section 2.3 describes the modelling of the quantum memory's interaction with a finite temperature environment as well as how it can be numerically simulated. Then the details of the error decoding procedure we employ are outlined in section 2.4.

2.1. Stabiliser codes

This section is a brief review of the construction of quantum stabiliser codes using the toric code as a particular example. This is purely to provide the necessary framework for the work discussed in the following chapter. Many excellent introductions and in depth reviews are available in the literature such as [16, 17].

2.1.1. Construction

Error correcting codes provide a way to store information robustly by using a large number of physical qubits (the fundamental unit of quantum data) to encode a small number of logical qubits. This is achieved by specifying a particular subspace of a much larger Hilbert space as the logical degrees of freedom. The remaining degrees of freedom are used to detect and attempt to correct errors. Stabiliser codes do this in an elegant way by specifying a set of mutually commuting operators whose joint +1 eigenspace define the

logical subspace, called the *code space*. This gives them an efficient description purely in terms of these operators.

A stabiliser code is defined over n qubits. Quantum states of these qubits live in a Hilbert space of dimension 2^n. The code space is defined within the full Hilbert space by fixing the eigenvalues of m independent, mutually commuting operators. These operators are called *stabilisers* S_r where $r = 1, \ldots, m$ and the code space is defined by

$$S_r \left|\Psi\right\rangle = +1 \left|\Psi\right\rangle \text{ for all } r.$$

The stabilisers may all be chosen to be hermitian and so correspond to some set of observables. Fixing these stabiliser degrees of freedom restricts the Hilbert space dimension of the code space to $2^{n-m} \equiv 2^k$ and so we encode k logical qubits.

In the case of the toric code the physical qubits are arranged on the links of a periodic $L \times L$ square lattice, fig. 2.1, so that $n = 2L^2$. Embedding this two-dimensional lattice in three-dimensional space, the qubits live on the surface of a torus. The toric code has two kinds of stabiliser operators: plaquettes operators B_p that apply a Z to all qubits living on the links around a plaquette, and star operators A_v that act with an X on the links adjacent to a vertex. Examples of these operators are shown in fig. 2.1(a). Here we have used the convention of denoting the Pauli matrices with X, Y and Z, explicitly

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The independent, commuting stabiliser operators are the generators of the Abelian *stabiliser group* $\mathcal{S} = \langle S_1, \ldots, S_m \rangle$ where $-\mathbb{1} \notin \mathcal{S}$, meaning the group $\mathcal{S}$ is composed of all possible products of these generators S_r. The stabiliser group is a subgroup of the full *Pauli group* of the n qubits, $\mathcal{P}_n = \langle i\mathbb{1}, X_1, Z_1, \ldots, X_n, Z_n \rangle$. Here we have introduced the notation that X_r means Pauli X acting on qubit r, explicitly this is

$$X_r = \underbrace{\mathbb{1} \otimes \mathbb{1} \otimes \cdots \otimes \mathbb{1}}_{r-1} \otimes X \otimes \underbrace{\mathbb{1} \otimes \mathbb{1} \otimes \cdots \otimes \mathbb{1}}_{n-r}.$$

The internal (logical) degrees of freedom of the code space are manipulated by *logical operators*. These operators must not disturb the stabilisers, so commute with all elements

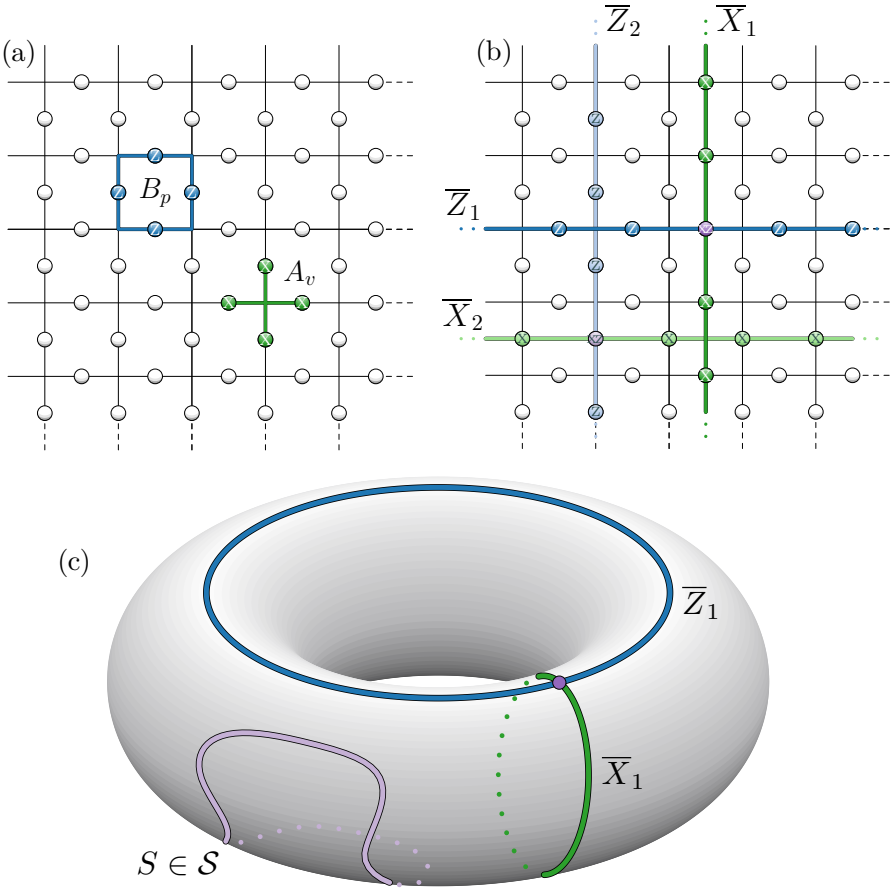


Figure 2.1. – Building blocks of the toric code. (a) examples of the plaquette B_p and star A_v stabiliser operators. Each plaquette p and vertex v, respectively, has a B_p and A_v operator associated with it. (b) the two anti-commuting pairs of logical operators of the toric code $\{\overline{X}_1, \overline{Z}_1\}$ and $\{\overline{X}_2, \overline{Z}_2\}$. (c) an element of the stabiliser group $S \in \mathcal{S}$ and one pair of logical operators $\{\overline{X}_1, \overline{Z}_1\}$ are drawn on a torus embedded in three-dimensional space to illustrate their relationships to topology. The logical operators run along different non-trivial cycles and only intersect at a single point. In contrast, all elements of the stabiliser group form trivial cycles on the surface of the torus.

of $\mathcal{S}$. However, they cannot be part of the stabiliser group or they would act trivially on encoded states. Logical operators can always be chosen as anti-commuting pairs of observables from the Pauli group $(\overline{X}_l, \overline{Z}_l) \in \mathcal{P}_n$ for $l = 1, \ldots, k$ where $\overline{X}_l\,\overline{Z}_l = -\overline{Z}_l\,\overline{X}_l$. Additionally, logical operators from separate pairs will commute, that is $\overline{X}_r\,\overline{Z}_s = \overline{Z}_r\,\overline{X}_s$ when $r \neq s$. The logical operators are only defined up to their product with the stabiliser group.

In the case of the toric code these conditions on the logical operators, that they must commute with all the stabilisers but cannot be generated from a product of them, has a topological interpretation. The logical operators of the toric code are strings of X or Z's that run around *non-trivial* cycles of the torus. Non-trivial cycles of the torus are loops that cannot be shrunk to a point without cutting and gluing them closed again, intuitively they are loops that pass around or through the 'hole' of the donut. The toric code encodes two logical qubits. The logical operators $\{\overline{X}_1, \overline{Z}_1\}$ and $\{\overline{X}_2, \overline{Z}_2\}$ are illustrated in fig. 2.1(b). To make the construction of these operators clearer, the $\overline{Z}$ logical operators are strings on the edges of the lattice and each correspond to one of the two different non-trivial cycles. The $\overline{X}$ logical operators that are conjugate to them are strings on the edges of the *dual graph* (i.e. the lattice obtained by replacing plaquettes with vertices and vice versa) and are associated with the opposite non-trivial cycle to their $\overline{Z}$ partner. In this way the $\overline{X}_i$ and $\overline{Z}_i$ operators only meet at a single qubit and so anti-commute with each other, whereas $\overline{Z}_1$ does not act on any of the same quibts as $\overline{X}_2$ or $\overline{Z}_2$ and so they commute (with the same being true for $\overline{X}_1$). In contrast to the logical operators, the stabiliser group (products of the A_v and B_p operators) only contains strings that are boundaries of regions of the lattice. In particular this includes strings with trivial topology, meaning they can be shrunk to a point without cutting the string. It additionally includes even numbers of non-trivial strings, which form the boundaries of 'strip' regions that are periodic in one direction. Recall that logical operators are only defined up to products with the stabiliser group, hence the $\overline{Z}$ and $\overline{X}$ operators defined above can be deformed with these loops and still remain logical operators. One pair of logical operators and an example element of the stabiliser group are drawn on a three-dimensional torus in fig. 2.1(c) to help illustrate the role of topology.

2.1.2. Detecting and correcting errors

Modelling arbitrary errors acting on a quantum code is in general a hard problem. A simplified approach that is often employed is to study *Pauli error channels*. These channels act on a state of the code ρ through the error map

$$\mathcal{N}(\rho) = (1-p)\,\mathbb{1} + \sum_i p_i E_i\,\rho\,E_i \quad \text{where} \quad p = \sum_i p_i \tag{2.1}$$

where $E_i \in \mathcal{P}_n$ is the complete Pauli group of the n physical qubits. Errors such as these act incoherently. They transform the state into a probabilistic mixture, described by distribution p_i, where each component suffers the Pauli error E_i. Applying Pauli errors E_i to stabiliser states yields another stabiliser state, only with a different set of eigenvalues, and thus the action of Pauli error channels on stabiliser codes can be efficiently simulated. In contrast, arbitrary unitary operators acting on the code would also include coherent components, which would enter eq. (2.1) in the form of terms like $(E_i\,\rho - \rho\,E_i)$. These cannot be efficiently simulated in general as their action would move the state out of the stabiliser eigenbasis. This is problematic as errors such as these are likely to be common in real quantum devices. Fully analysing the effect of coherent errors remains an open problem and an active area of research [64–66]. However, recent progress suggests that for large enough topological codes locally acting coherent physical errors result in an effective error channel acting on the logical subspace that is incoherent [64]. This result, and others, suggests that the simplified modelling of Pauli error channels is nevertheless an effective way to analyse the robustness of quantum codes.

We will work only in the framework of Pauli error channels. Every member of the Pauli group either commutes or anti-commutes with each stabiliser S_r; therefore, if a non-trivial error occurs its signature will be that some of the stabiliser eigenvalues will be flipped to -1. The measurement outcomes of the stabilisers following an error E are called the *error syndrome*. This syndrome is not unique as multiplying the error by any element of the stabiliser group $E' = E\,S$, $S \in \mathcal{S}$, will not change the syndrome.

Generic members of the Pauli group can be decomposed (up to phases) into strings of operators acting on the toric code. Closed strings belong to either the stabiliser group or the set of logical operators. Both of these commute with all the stabilisers so have

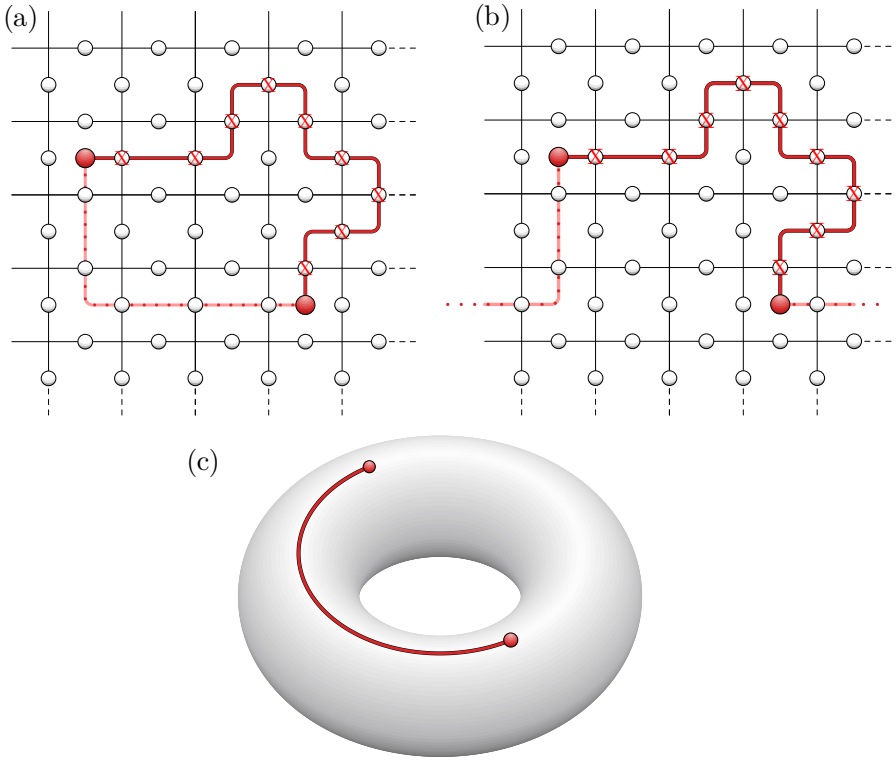


Figure 2.2. – An open chain of errors on the toric code creates two non-trivial syndromes by flipping the eigenvalue of the stabilisers at each ends. In this case the flipped stabilisers are the B_p plaquette operators indicated with red spheres.. Two proposed correction operators result in either: (a) no damage to the encoded state or (b) a logical error. In (a) and (b) the error that has occurred is indicated with a solid red line as well as X operators on the qubits, and the proposed correction operators are drawn as light/dashed lines. (c) a worst case error configuration for the toric code. A decoder cannot reliably correct such an error.

trivial syndromes. Open strings, on the other hand, will anti-commute with the stabilisers at the end points of the string and so generate a non-trivial error syndrome that will be detected by measuring the stabilisers. Figure 2.2 (a) shows an open string of X's causing two flipped stabilisers, drawn as red balls, at the ends of the string.

We can try to correct an error by acting with a correction operator C. This operator will be such that the product of the initial error E with C has trivial syndrome. The correction succeeds if $C\,E \in \mathcal{S}$, meaning the final action on the state is just multiplication with a stabiliser, and fails if $C\,E = \overline{L}$ for some logical operator $\overline{L}$. The job of choosing a correction operator C is carried out by a *decoder*. One way to evaluate the strength of a code is to consider the smallest worst case error. This is the error (or errors) that act on the smallest number of qubits necessary to produce a syndrome pattern that is impossible to reliably decode. Considering errors like these leads to the concept of *code distance*. The

code distance, d, is defined as the *weight* (the number of qubits acted on non-trivially) of the smallest logical operator. Since the logical operators are only defined up to products taken with the stabilisers, the code distance is obtained by minimising the weight of the logical operators with respect to the action of the stabiliser group. The smallest worst case error will then be an error that acts on $^{d}/_{2}$ qubits.

Toric code correction operators will be open string Pauli operators that produce the same syndrome as the error. If the combined closed string (error plus correction) has trivial topology the final operation applied belongs to the stabiliser group and the error has successfully been corrected, fig. 2.2(a). However, if the total loop wraps around a non-trivial cycle a logical error has occurred, fig. 2.2(b). For the toric code defined on an $L \times L$ square lattice all minimum weight logical operators have the same weight of L (these are the logical operators as drawn in fig. 2.1), hence the distance of the code is $d = L$. The least-weight worst case errors are strings of length $^{L}/_{2}$ that wind halfway around a non-trivial cycle, such as the error configuration drawn in fig. 2.2(c). Errors like these provide no information to a decoder about what direction the error string lies in and cannot be reliably corrected.

In principle it is not necessary to actually carry out the correction operation. If we can decide through decoding that the correction we would like to apply is C then we can store this information and instead use it to adapt future elements of the computation in such a way that the error is dealt with [16]. This correction operator C is then called the *Pauli frame* of the encoded state. Depending on the situation it will take different amounts of effort to account for a non-trivial Pauli frame. For complex computations (involving non-Clifford gates) this could mean having to alter the operators that are applied at later steps of the computation. Whereas, in the simplest case it could mean we just have to reinterpret the result of final measurement outcomes.

2.2. Quantum memory Hamiltonians

Quantum memories use energetic protection to try and suppress the development of problematic errors. Information is stored in the degenerate ground state eigenspace of a many-body (N qubit) Hamiltonian, $H_{\text{mem}} = -\Delta \sum_k P_k$ where $P_k \in \mathcal{P}_N$, such that errors have

an energy cost and so are suppressed. This general setup is greatly simplified by restricting the Hamiltonian terms to a stabiliser group over the qubits $P_k \in \mathcal{S}$. In this case H_{mem} is exactly solvable and we can use the language of error correcting codes to understand many of its properties. The ground state eigenspace of H_{mem} is the code space of the stabiliser code $\mathcal{S}$ and the degeneracy of the ground state is resolved by the logical operators of the code. Excitations are described by the eigenvalues of the stabiliser generators.

For the toric code the simplest such Hamiltonian is to use the stabiliser generators as the interaction terms,

$$H_{\text{tc}} = -\frac{\Delta_p}{2}\sum_p B_p - \frac{\Delta_v}{2}\sum_v A_v \quad \text{where} \quad \Delta_{p,v} > 0. \tag{2.2}$$

It's easily seen that the code states, for which all $\langle B_p \rangle = \langle A_s \rangle = 1$, minimises its energy; therefore the toric code codespace is the ground state eigenspace. We can construct projection operators onto the (-1) flipped eigenvalue of the stabilisers

$$n_p = \frac{\mathbb{1} - B_p}{2} \qquad \& \qquad n_v = \frac{\mathbb{1} - A_v}{2}$$

and, up to non-physical constant terms, rewrite the Hamiltonian using these projectors

$$H_{\text{tc}} = \Delta_p \sum_p n_p + \Delta_v \sum_v n_v \,. \tag{2.3}$$

This translates flipped stabilisers into quasiparticle excitations of two species n_p and n_v. In this view, the ground state subspace is the vacuum of these particles. Error operators create and transport the n_p and n_v. As these particles live at the end of strings, both n_p and n_v particles always exist in pairs.

2.2.1. Topological order and toric code anyons

The toric code Hamiltonian is a topologically ordered stabiliser Hamiltonian. A colloquial definition of topological order is that there is no local order parameter to resolve the degenerate ground states. We already know that the only operators that lift the degeneracy of the ground states of the toric code are the non-trivial loop operators, the size of which

scales with the dimensions of the system. Topological order ensures the ground state degeneracy is robust to perturbations, such as stray magnetic fields [17]. Specifically, it has been shown that for topologically ordered stabiliser Hamiltonians the gap induced in the ground state subspace by weak local perturbations is exponentially small in the system size [67]. If this was not the case such perturbations would lead to coherent *dephasing* errors of the logical states, as different states in the code space would acquire different phases under the Hamiltonian evolution of the system. This resilience is an essential characteristic if the system is to serve as a quantum memory. However, this picture of coherent errors is far from complete and misses some important aspects. For example this neglects code preparation errors, when the state is not a perfect eigenstate of the Hamiltonian, as well as the possibility that the perturbations induces dynamics in the quasiparticle excitations that is damaging to the encoded information. For an extended discussion of these topics see section II.F of [17]. Nevertheless, topologically ordered systems provide the best known examples of physical systems whose natural properties help to suppress quantum errors.

The quasiparticles of the toric code introduced above, n_p and n_v, are Abelian anyons. Specifically, they belong to the *quantum double* of the Abelian group $\mathbb{Z}_2$ [35]. The quasiparticles are conventionally labelled e and m. This anyonic picture gives another view of the logical states. Any string operator can be viewed as the trajectory of a pair of e or m particles moving over the lattice. Therefore, logical operators correspond to particles winding around a non-trivial cycle of the torus before finding each other and annihilating. This annihilation returns the system to its ground state but leaves its particle flux. In the lattice model of the toric code this particle flux corresponds to a closed string with non-trivial topology (the logical operator), which has generated the motion of the particle. Measurement of the conjugate logical operator can detect the presence or absence of this flux loop. A physical picture of the logical degrees of freedom is that they are the occupations of these winding fluxes.

2.3. Modelling finite temperatures

We are interested in errors arising from the interaction between the quantum memory and a thermal environment. The thermal environment is modelled by a large auxiliary quantum system. The initial state of this auxiliary system is a thermal state at temperature T and we call it a *thermal bath*. The interaction between the memory and the bath causes the two systems to become entangled. Thus, over time it becomes harder to reliably extract the information initially stored in the memory when we can only perform measurements on the memory's degrees of freedom.

In order to have a tractable model of the interaction between the memory and the bath we assume the coupling is weak and that the bath is much larger than the memory. Under these assumptions, the state of the bath effectively does not change due to the interaction and it can be considered to remain in a thermal state at all times. For all practical purposes, the information transferred from the memory into entanglement is now lost all together. This is the *Markovian* limit. For a thermal bath in this limit the memory evolves towards its Gibbs state

$$\rho_\beta = \sum_j \frac{e^{-\beta\varepsilon_j}}{\mathcal{Z}} |e_j\rangle\langle e_j| \ , \tag{2.4}$$

where $|e_j\rangle$ are some complete set of energy eigenstates of the memory with energies ε_j, $\beta = 1/T$ is the inverse temperature of the bath (with Boltzmann's constant set equal to one) and $\mathcal{Z} = \sum_j e^{-\beta\varepsilon_j}$ is the canonical partition function.

An additional assumption gives a particularly simple model of the dynamics. If we assume that the memory-bath interactions are local and independent then the memory dynamics can effectively be described by a rate equation. We can visualise such a bath as injecting or extracting lumps of energy from the memory so causing jumps between energy eigenstates. Each jump occurs with some characteristic rate that depends only on the energy difference. If V is a process jumping the memory from energy ε_i to ε_f, so causing the energy of the bath to change by $\omega_V = -(\varepsilon_f - \varepsilon_i)$, then it occurs with frequency

$$\gamma(\omega_V) = \frac{\kappa\,\omega_V}{1 - e^{-\beta\omega_V}} \ , \tag{2.5}$$

where κ has units of $1/(\text{energy} \times \text{time})$ and sets the timescale of energy exchange between

the memory and the environment. These rates can be seen to satisfy the property of detailed balance, $\gamma(\omega_V) = e^{\beta\omega_V}\gamma(-\omega_V)$.

2.3.1. Deriving the rate equation

We will give the derivation of the rate equation, eq. (2.5), from a weak coupling of the memory to a thermal bath [68]. Very general and powerful mathematical frameworks have been developed in order to derive weak coupling master equation [69]. These methods provide a more rigorous route to obtaining eq. (2.14). Rather than being rigorous, the aim of this derivation is to highlight the physical origins of the different parts of the final result as well as the assumptions involved in reaching it.

Rather than being rigorous, the aim of this derivation is to highlight the physical origins of the different parts of the final result as well as the assumptions involved in reaching it.

The memory, with Hamiltonian H_{mem}, interacts with a bath described by H_{bath}. The initial state is an uncorrelated product state between a pure state of the memory and the thermal state of the bath, $\rho_{\text{tot}} = |\Psi\rangle\langle\Psi| \otimes \rho_{\text{th}}$. Taking a generic form of the interaction gives a total Hamiltonian

$$H \;=\; H_{\text{mem}} \otimes \mathbb{1}_{\text{bath}} + \mathbb{1}_{sys} \otimes H_{\text{bath}} + \sum_{\alpha} W_\alpha \otimes f_\alpha \;\equiv\; H_M + H_B + H_I \,,$$

where W_α and f_α are assumed to be hermitian and act on the memory and bath respectively. The interaction operators on the memory W_α are *ergodic*, meaning that the interaction can address all states of the memory (there are no *dark states* that decouple). To describe the dynamics we move to the interaction picture. Interaction picture operators are denoted with a tilde and defined by $\tilde{O} = e^{i(H_M+H_B)t}\, O\, e^{-i(H_M+H_B)t}$. The state $\tilde{\rho}_{\text{tot}}$ of the coupled memory and bath evolves according to

$$\frac{\partial}{\partial t}\tilde{\rho}_{\text{tot}} = -i[\tilde{H}_I(t), \tilde{\rho}_{\text{tot}}(t)] \,. \tag{2.6}$$

A formal solution to eq. (2.6) is obtained by integrating both sides. This formal solution

can then be substituted back into the right hand side of eq. (2.6) to obtain

$$\frac{\partial}{\partial t}\tilde{\rho}_{\text{tot}} = -i[\tilde{H}_I(t), \tilde{\rho}_{\text{tot}}(0)] - \int_0^t \mathrm{d}s\, [\tilde{H}_I(t), [\tilde{H}_I(s), \tilde{\rho}_{\text{tot}}(s)]]\,. \tag{2.7}$$

We are only interested in the dynamics of the memory so we next trace out the bath to find the evolution of the reduced state $\tilde{\rho} = \text{tr}_{\text{bath}}(\tilde{\rho}_{\text{tot}})$. The partial trace of the first term on the right hand side can be chosen to vanish, $\text{tr}_{\text{bath}}([\tilde{H}_I(t), \tilde{\rho}_{\text{tot}}(0)]) = 0$. In particular this will be the case since the initial state of the bath, the thermal state, is diagonal in its energy eigenbasis. Whereas the interaction operator acting on the bath is off-diagonal, since it encodes energy exchanges between the memory and the bath. The derivation so far is exact, we now begin to simplify eq. (2.7) with approximations.

In the weak coupling limit we make the *Born approximation*. This replaces $\tilde{\rho}_{\text{tot}}(s)$ on the right hand side of eq. (2.7) with the factorised $\tilde{\rho}(s) \otimes \rho_{\text{th}}$ at all times. The reduced dynamics of the memory are then described by

$$\frac{\partial}{\partial t}\tilde{\rho} = -\int_0^t \mathrm{d}s\, \text{tr}_{\text{bath}}\Big([\tilde{H}_I(t), [\tilde{H}_I(s), \tilde{\rho}(s) \otimes \rho_{\text{th}}]]\Big)\,. \tag{2.8}$$

One possible justification of this is that, since the coupling is weak, the state of the bath effectively does not change. Or if it does, then since it is so much larger than the memory it relaxes back to its thermal state on timescales that are so short the change can be neglected. However, a more sophisticated analysis can show that this does not need to be the case [70]. Even in the weak coupling limit the state of the bath might change but eq. (2.8) remains the correct description of the reduced dynamics of the memory.

If we substitute the expression $\tilde{H}_I = \sum_\alpha \tilde{W}_\alpha \otimes \tilde{f}_\alpha$ into eq. (2.8) and expand the double commutator we find it is possible to use the cyclic property of the trace to refactorise the integrand into

$$\begin{aligned}\frac{\partial}{\partial t}\tilde{\rho} = -\sum_{\alpha,\beta}\int_0^t \mathrm{d}s\, \Big(&[\widetilde{W}_\alpha(t), \widetilde{W}_\beta(s)\tilde{\rho}(s)]\, \text{tr}_{\text{bath}}(\tilde{f}_\alpha(t)\tilde{f}_\beta(s)\rho_{\text{th}}) \\ &-[\widetilde{W}_\alpha(t), \tilde{\rho}(s)\widetilde{W}_\beta(s)]\, \text{tr}_{\text{bath}}(\tilde{f}_\beta(s)\tilde{f}_\alpha(t)\rho_{\text{th}})\Big)\,.\end{aligned} \tag{2.9}$$

The correlation functions $C_{ab}(t_1, t_2) = \text{tr}_{\text{bath}}(\tilde{f}_a(t_1)\tilde{f}_b(t_2)\rho_{\text{th}})$ now encapsulate all remain-

ing reference to the bath. Since $\rho_{\rm th}$ is a stationary state, these correlations only depend on the time difference, $C_{ab}(t_1, t_2) = C_{ab}(t_1 - t_2)$. To continue, we expand the operators $\widetilde{W}_\alpha$ in terms of their Fourier components. Using the definition of the interaction picture operators, we find these are given by

$$\widetilde{W}_\alpha(t) = e^{iH_{\rm mem}t}\, W_\alpha\, e^{-iH_{\rm mem}t} = \sum_\omega V_\alpha(\omega)\, e^{i\omega t}\,.$$

The set of ω are energy differences between different eigenstates of $H_{\rm mem}$ and the operators $V_\alpha(\omega)$ are the parts of W_α that couple those two different energies. We note that since the W_α are hermitian this means $V_\alpha(-\omega) = V_\alpha^\dagger(\omega)$. Substituting into eq. (2.9) and changing the integration variable to $\tau = t - s$ gives

$$\frac{\partial}{\partial t}\tilde{\rho} = -\sum_{\alpha,\beta}\sum_{\omega,\nu}\int_0^t \mathrm{d}\tau\, e^{i(\omega+\nu)t}\, e^{-i\nu\tau}\Big([V_\alpha(\omega), V_\beta(\nu)\tilde{\rho}(t-\tau)]\, C_{\alpha\beta}(\tau) \\ - [V_\alpha(\omega), \tilde{\rho}(t-\tau)V_\beta(\nu)]\, C_{\beta\alpha}(-\tau)\Big)\,. \tag{2.10}$$

We now make a *rotating wave approximation* and say that the only terms that will contribute to the sum over ω and ν are those where ν is equal to $-\omega$. Other terms will average out due to the fast oscillating $e^{i(\omega+\nu)t}$.

The time taken for the bath to relax back to equilibrium is much shorter than any timescale of the memory. We therefore expect the bath correlation functions to be peaked around $\tau = 0$ and decay very fast. While, in the same amount of time we expect the state of the memory to hardly have changed. This justifies replacing $\tilde{\rho}(t-\tau)$ inside the integral with $\tilde{\rho}(t)$ and extending the upper bound of the integral to infinity. This is the *Markovian approximation*, since in this regime the future evolution of the memory is determined only by its current state and not its history. The equation of motion is now

$$\frac{\partial}{\partial t}\tilde{\rho} = -\sum_{\alpha,\beta}\sum_{\omega}\Bigg([V_\alpha(\omega), V_\beta(-\omega)\tilde{\rho}]\int_0^\infty \mathrm{d}\tau e^{i\omega\tau} C_{\alpha\beta}(\tau) \\ -[V_\alpha(\omega), \rho\, V_\beta(-\omega)]\int_0^\infty \mathrm{d}\tau e^{i\omega\tau} C_{\beta\alpha}(-\tau)\Bigg)\,. \tag{2.11}$$

To simplify further we define the matrix $\Gamma_{\alpha\beta}(\omega) = \int_0^\infty \mathrm{d}\tau e^{i\omega\tau} C_{\alpha\beta}(\tau)$, which is equal to the first integral. Using the fact that $C_{\beta\alpha}(-\tau) = C^*_{\alpha\beta}(\tau)$, the second integral is given

by $\Gamma^*_{\alpha\beta}(-\omega)$. We break $\Gamma_{\alpha\beta}$ into its hermitian and anti-hermitian parts and write it as $\Gamma_{\alpha\beta}(\omega) = {}^1\!/_2\, g_{\alpha\beta}(\omega) + {}^i\!/_2\, h_{\alpha\beta}(\omega)$ where

$$g_{\alpha\beta}(\omega) = \Gamma_{\alpha\beta}(\omega) + \Gamma^*_{\alpha\beta}(\omega) = \int_{-\infty}^{\infty} \mathrm{d}\tau e^{i\omega\tau} C_{\alpha\beta}(\tau)\,,$$
$$h_{\alpha\beta}(\omega) = {}^1\!/_i\big(\Gamma_{\alpha\beta}(\omega) - \Gamma^*_{\alpha\beta}(\omega)\big)\,.$$

Note that the hermitian part $g_{\alpha\beta}(\omega)$ is the Fourier transform of the bath correlation functions. Substituting these expressions into eq. (2.11) and manipulating eventually gives

$$\frac{\partial}{\partial t}\tilde{\rho} = \frac{1}{2}\sum_{\alpha,\beta}\sum_{\omega}\Big(-i\, h_{\alpha\beta}(\omega)\,[V_\alpha(\omega)V_\beta(-\omega), \tilde{\rho}] + g_{\alpha\beta}(\omega)\big(2\,V_\beta(-\omega)\tilde{\rho}V_\alpha(\omega) - \{V_\alpha(\omega)V_\beta(-\omega), \tilde{\rho}\}\big)\Big)\,. \tag{2.12}$$

The first term, typically called the shift Hamiltonian, renormalises the energy levels of the memory. It has been shown that this term cannot speed up the decoherence of the memory [21, 24] and so for simplicity we will neglect it when studying the dynamics.

We now start to make choices about the type of bath the memory is coupled to and the details of the interaction. We independently couple each physical qubit to a bath made up of an infinite number of harmonic oscillators. An ergodic choice of W_α that act on single qubits are Pauli X's and Z's. We additionally assume that the Pauli X and Z thermal errors on a single qubit are independent. The independent coupling means there are no correlations between different baths, $g_{\alpha\beta}(\omega) = g_\alpha(\omega)\,\delta_{\alpha\beta}$. Finally, returning to the Schrodinger picture, we arrive at the Markovian master equation for the reduced dynamics of the memory

$$\dot{\rho} = -i[\rho, H_{\mathrm{mem}}] + \mathcal{L}(\rho)\,, \tag{2.13}$$

where the Liouvillian super-operator $\mathcal{L}$ encodes the effect of the environment on the memory and is given by

$$\mathcal{L}(\rho) = \sum_{\alpha}\sum_{\omega} g_\alpha(\omega)\left(V_\alpha(\omega)^\dagger \rho V_\alpha(\omega) - \frac{1}{2}\{V_\alpha(\omega)^\dagger V_\alpha(\omega), \rho\}\right) \tag{2.14}$$

The power spectrum of the bath $g_\alpha(\omega)$ describes how quickly or slowly the bath can exchange energy with the memory. It must satisfy detailed balance $g_\alpha(\omega) = e^{-\beta\omega}\, g_\alpha(-\omega)$

in order to ensure the thermal state of the memory is the equilibrium state.

The initial state of the memory is a pure ground state. This is the case we will often consider in the quantum memory setting, however all of the following also applies to states that are mixed in the energy eigenbasis of the memory for example probabilistic mixtures of different ground states. If there is no initial coherence in the energy eigenbasis the dynamics described by eq. (2.13) cannot introduce coherence between energy eigenstates [21]. The density matrix of the memory will therefore be diagonal in the energy eigenbasis at all times in the evolution $\rho(t) = \sum_i p_i(t) \, |e_i\rangle\langle e_i|$. Substituting this form of ρ into the master equation we can show that it takes on a classical form

$$\dot{p}_k = \sum_{j \neq k} \left(\Gamma_{j \to k} \, p_j - \Gamma_{k \to j} \, p_k \right) \tag{2.15}$$

where $\Gamma_{j \to k}$ is the rate of transition from state j to k. If we assume the $g_\alpha(\omega)$ satisfy detailed balance we find the following expression for the rates. Let V be a process that causes the transition $j \to k$ with change in the energy of the bath $\omega_V = -(\varepsilon_k - \varepsilon_j)$, the rates are then written

$$\Gamma_{j \to k} = \Gamma(\omega_V) = \frac{g(\omega_V)}{|1 - e^{-\beta \omega_V}|} \, . \tag{2.16}$$

With the final assumption that the bath has Ohmic spectral density and large cut-off energy, $g(\omega_V) \propto \omega_V$, eq. (2.16) yields the rate equation, eq. (2.5). This choice of bath power spectrum is motivated by the spin-boson model [71].

If we were studying the decoherence of a specific physical architecture for a quantum memory (such as trapped ions or NV centres in diamonds) we could not neglect the constant prefactor of $g(\omega_V) = \kappa \, \omega_V$. This constant κ sets the fundamental timescale for interactions between the memory and bath in physical units. However, when studying quantum memories from a quantum coding perspective we are motivated by the more abstract question of how the *coherence time* scales with temperature and system size.

2.3.2. Coherence time

An error correcting scheme can only really be evaluated in the context of some particular error model and different error models will have different figures of merit. In the case of

thermal errors evolving under some master equation this is the coherence time.

The coherence time τ is defined differently in different studies. However, the principle is that τ quantifies the time up to which we expect to be able to reliably recover stored information from the memory using a decoder. Some possible definitions are that τ is the expectation value of the time the decoder first fails, or that τ is the time that the decoder success rate first falls below some threshold such as 99%.

2.3.3. Simulating thermal physics

Under the simplifying assumptions made above, the thermal decoherence of the memory is described by a classical master equation, eq. (2.15), with transition rates that only depend on energy differences, eq. (2.5). To study the effect on encoded information, the time evolution of an initial pure ground state under these dynamics can be approximated with Monte Carlo methods [72, 73]. The rest of this section outlines the steps of the evolution algorithm we employ.

At time $t = 0$ the memory is in an energy (and logical) eigenstate, $|\Psi_0\rangle$. As time progresses, thermal processes jump the system to other energy eigenstates by applying single spin X or Z operators. These jumps occur probabilistically producing a mixed state $\rho(t)$ for times $t > 0$. We numerically sample from the ensemble $\rho(t)$ by averaging over many simulated sequences of discrete jumps between pure states $|\Psi_0\rangle \to V_1 |\Psi_0\rangle \to V_2 V_1 |\Psi_0\rangle \to \ldots$, where the processes V_1, V_2, ... and the times they occur t_1, t_2, ... are chosen randomly with appropriate probabilities.

After i steps the simulation time is t_i and the memory is in state $|\Psi_i\rangle$, which has energy ε_i. For all possible thermal processes that can occur V, we evaluate their rates $\gamma(\omega_V)$ with respect to ε_i and ε_f, which is the energy of $V |\Psi_i\rangle$. The time we would expect to wait for each V to occur is exponentially distributed with rate $\gamma(\omega_V)$. Therefore, the time we expect to wait for any process to occur is also exponentially distributed with total rate $R = \sum_V \gamma(\omega_V)$. The time until the next jump $t_{i+1} = t_i + \delta t$ can then be chosen by drawing from this exponential distribution using the formula

$$\delta t = -\frac{\ln(\mathtt{r})}{R}\,, \tag{2.17}$$

where `r` is a random number chosen uniformly from the interval $(0, 1]$.

The relative likelihoods of each possible V occurring at time t_{i+1} are also given by the rates. The probability that V will be the next process to take place is

$$p_V = \frac{\gamma(\omega_V)}{R}\,. \tag{2.18}$$

Using these probabilities we randomly select the next thermal process V_{i+1}. This gives the state after the next jump $|\Psi_{i+1}\rangle = V_{i+1}\,|\Psi_i\rangle$. Unless the simulation has met some stop condition we now update the rates $\gamma(\omega_V)$ and repeat the above to continue to the next step. (As a practical note, since the toric code Hamiltonian is a sum of local commuting terms only a small number of the $\gamma(\omega_V)$ will have changed. We can identity these and only recompute the rates where necessary.)

This simulation scheme can easily be extended to simulate the evolution of initial mixed states of the memory. In this case the pure state chosen at the start of the simulation $|\Psi_0\rangle$ would be chosen at random by drawing from the probability distribution corresponding to the mixed state at $t = 0$.

2.4. Decoding errors

In active memories errors are handled through a repeated loop of measuring stabilisers, decoding errors and applying corrections. Though their application is more limited than in active schemes, decoders still play an important role in passive memories. An ideal passive memory would be left untouched while storing information. However, we would still want to decode when reading out the stored state. This final decoding step will always be necessary to remove residual errors, even for a truly self-correcting memory [17]. To reiterate a previous point, the coherence time τ is defined in relation to this decoding step. It quantifies how long we expect to be able to wait and still reliably decode. As such the particular choice of decoder will effect the value of τ, though we do not expect the general scaling of τ with temperature and system size to change.

The job of a decoder is to begin with an error syndrome pattern and choose a correction operator that will remove all error syndromes in a way that is least likely to cause logical

errors. For the toric code this means closing open strings into closed strings. The decoder succeeds if all the resulting closed strings have trivial homology, and it fails if the final operator has non-trivial homology. When decoding the toric code we can choose to only address Pauli X errors and their error syndromes of plaquette stabilisers. Pauli Z errors and the associated flipped star stabilisers can be dealt with independently using the same procedure.

There are many types of decoders developed for different situations [27, 74–77]. These can be specialised to particular codes or widely applicable to broad classes of codes. The decoder most often applied to the toric code is a specialised decoder, the minimum weight perfect matching algorithm [27]. However, we instead employ one of the most flexible types of decoder: a hard decision renormalisation group (HDRG) decoder, which is alternatively called a clustering decoder [25, 78].

2.4.1. Clustering decoder

A clustering decoder divides the syndromes into spatially local subsets, clusters, that can be corrected by local operators. With respect to constructing the clusters, local is taken to mean that there is some length z such that for every syndrome in the cluster there is at least one other member of the cluster that is at most distance z away. The cluster can be corrected locally if it is possible to remove all syndromes in the cluster with a correction operator that acts only on the physical qubits within a box bounded by the edges of the cluster. Such a cluster is called *neutral*. In the case of the toric code it is particularly easy to identify neutral clusters as it simply means they contain an even number of syndromes. If there are still un-neutralised clusters amongst the set of length z clusters the size is increased $z \to z + 1$ and the set of $z + 1$ clusters is considered. This is repeated until all the clusters are neutral. The choice of linearly growing clusters is one particular choice, it is also possible to consider exponentially growing clusters $z \to 2z$ [25] or other growth schemes. The algorithm is *greedy* in the sense that once all the clusters of length z have been constructed if any are found to be neutral the correction operator is applied straight away before constructing the $z + 1$ clusters. The syndromes in those clusters then disappear and are not included in the $z + 1$ clusters. (It is known that

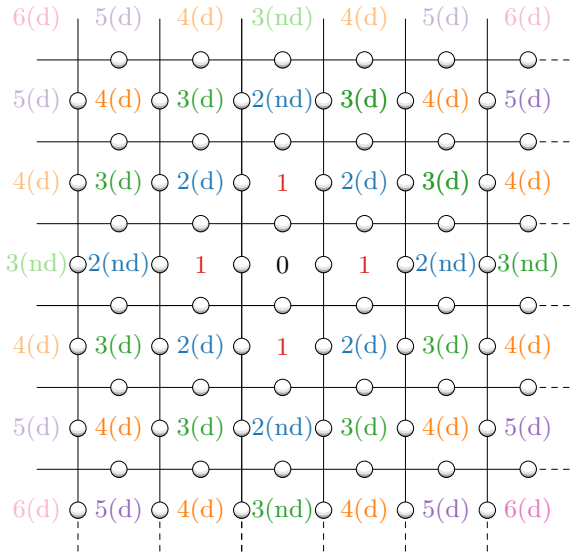


Figure 2.3. – Distances of all sites from a central site (site-0) labelled on a 6×6 toric code, based on the distance metric described in the text. The (d) or (nd) labels indicate sites linked to site-0 by degenerate or non-degenerate paths of a given length. Sites on the left/top edge of the lattice are faded out to indicate they are repeats of the same site labelled on the right/bottom edge due to the periodic boundaries. For distances 3(nd) and greater the finite size of the system effects the number of sites at each distance. In particular, there are only two distinct sites at distance 3(nd) rather than four. Additionally, there are no sites at distance 4(nd) or 5(nd). The most distant site is at distance 6(d) and there is only one such site.

relaxing this greediness can improve the success of the decoder [79].) Additional to the locality condition on the clusters, non-neutralised clusters cannot overlap spatially. For every member of each cluster, all the remaining syndromes within distance z must be a member of the same cluster. We will next give an example of applying this algorithm to a particular toric code error to make the steps of the procedure clearer. However, before that, in order to properly define the size of a cluster a distance metric is needed.

The distance z between plaquettes is taken to be the minimum number of operators needed to connect the two plaquettes by an error chain (this is the Manhattan distance on the square lattice). An additional complication is added as an entropic correction [77]. For $z \geq 2$ some pairs of plaquettes are connected by multiple chains of length z while others are only connected by a single path of length z. Entropic considerations mean the two sites connected by multiple paths are more likely to be linked together by an error chain of length z than the two with only a single path. These entropically favoured sites with degenerate paths of length z are considered closer together than sites with non-degenerate paths of the same length, in the sense that they are treated as two separate increments of

the algorithm. This distance metric is illustrated in fig. 2.3.

To help illustrate the steps of algorithm, consider the pattern of errors on the 5×5 toric code shown in fig. 2.4(a). The syndromes are drawn as well as the actual error strings, which are unknown to the decoder. The first step is an initialisation step. At $z = 0$ clusters are created for each syndrome containing only the syndrome itself, fig. 2.4(b). Increasing the search distance to $z = 1$, nearest neighbour syndromes are joined into clusters. This leaves three independent clusters, fig. 2.4(c), one of which is neutral. The neutral cluster is indicated in blue and the (in this case unique) correction operator is shown. The algorithm now proceeds separately through distances $z = 2$(degenerate) and then $z = 2$(non-degenerate) trying to expand the clusters. The two remaining clusters do not become joined at these distances. At $z = 3$(degenerate) the two remaining cluster join and become neutral, fig. 2.4(d). A possible choice of correction operator acting within the box bounding the cluster is also shown. The $z = 1$ neutral cluster is not shown here since as soon as it became neutral after $z = 1$ the correction operator was applied. Finally, fig. 2.4(e) shows the combined action of the initial error and the correction operators. In this case the decoder has failed. The combined operator decomposes into a non-trivial loop of X's in both the vertical and horizontal directions meaning a logical $\overline{X}$ has been applied to both stored qubits. Using X stabilisers to deform the error chain we can bring it into the shape of the logical operators as originally defined, fig. 2.4(f).

2.4.2. Testing for success

We will simulate the building up of errors in the memory and then some time later try to decode. Rather than explicitly construct correction operators and check the homology of the final operator we instead take the following approach to testing if the decoder has been successful.

A choice of logical operators is fixed, for example $\overline{Z}_1$ and $\overline{Z}_2$ are chosen to be the strings of Z's applied to the qubits shown in fig. 2.5. With these operators specified, we can find the Pauli frame of the logical qubits induced by the errors. In the example of fig. 2.5(a) none of the physical error strings intersect the logical operators and so the logical Pauli frame is just identities. This is compared to the logical Pauli frame deduced

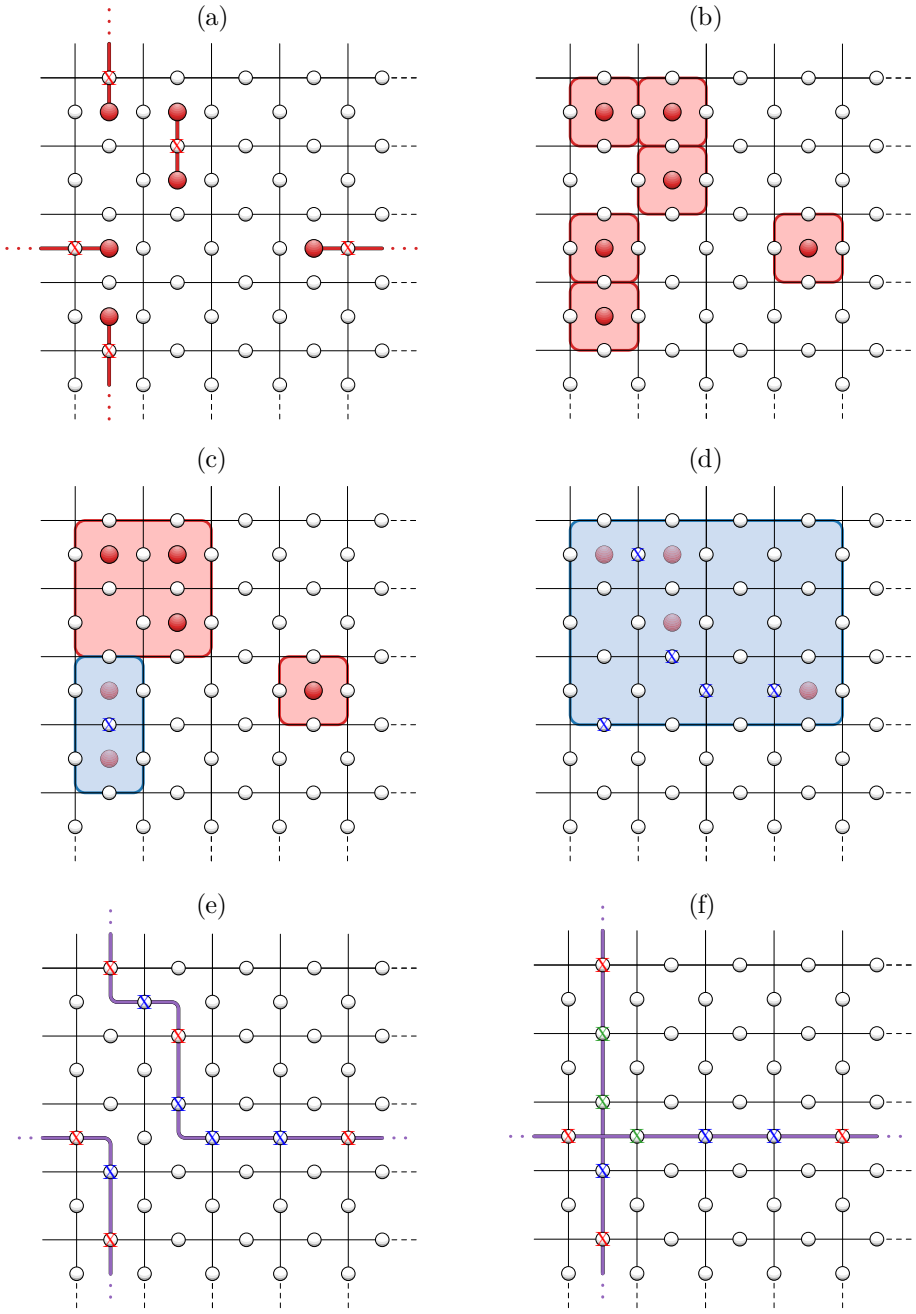


Figure 2.4. – Illustration of the working of the decoder. (a) The initial error pattern, the syndromes (red spheres) become the input for the decoder and the error strings are unknown. (b) Beginning of the algorithm, search distance $d = 0$. Trivial clusters are created (red boxes). (c) Search radius is expanded to $d = 1$ and all syndromes within that distance of each other are clustered. The lower left cluster is neutral and a correction operator that will remove these syndromes is shown. (d) When $d = 3$(degenerate) the remaining clusters combine and become neutral. A new correction operator is shown, in addition to the previously applied correction. (e) The total operator applied to the lattice. (f) Deforming the total operator with stabilisers to recover the two X logical operators shown in fig. 2.1. This means the decoder has failed and a logical error has occurred.

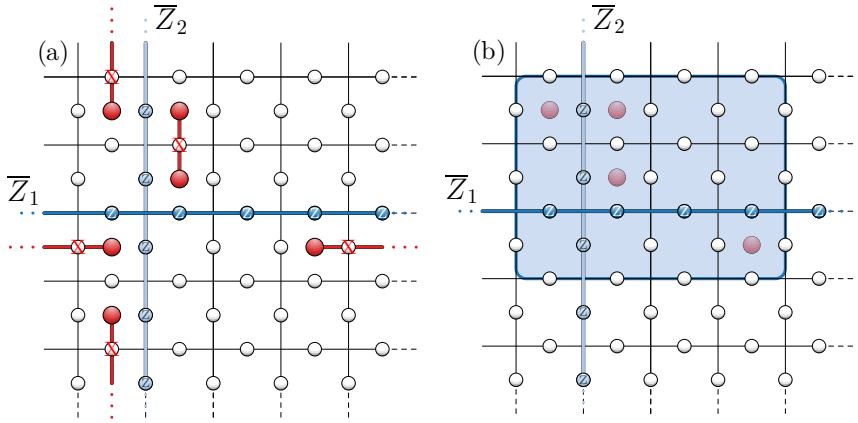


Figure 2.5. – Assessing decoder success by comparing the actual Pauli frame of the logical qubits to the Pauli frame deduced by the decoder. (a) based on the fixed choice of logical strings $\overline{Z}_1$ and $\overline{Z}_2$ the actual Pauli frame of the logical operators is trivial. (b) the result of decoding means that X strings will have to cross $\overline{Z}_1$ and $\overline{Z}_2$ to remove the cluster highlighted. Thus giving a deduced Pauli frame that has bit-flips on both logical qubits.

by the decoder. At the end of decoding we identify the clusters that intersect with our choice of logical operators. This is only the large cluster shown in fig. 2.5(b). For each logical operator we compute the *particle flux* crossing the logical string. In this case, for both $\overline{Z}_1$ and $\overline{Z}_2$, there is an odd number of syndromes on each side of the logical strings. This means that any correction operator will need to transport an anyon across $\overline{Z}_1$ and across $\overline{Z}_2$. Therefore, the deduced Pauli frame is that a bit-flip has occurred to both logical qubits.

If the future results of computations are adjusted according to our guess of a bit-flip Pauli frame we will introduce a logical error into our computation. This again means that the decoder has failed. Since this method only requires us to count the number of syndromes on each side of fixed logical strings it is a highly efficient way to assess decoder success.

Finally, before progressing to consider thermal errors we benchmark our implementation of the decoder against independent identically distributed (i.i.d.) errors to show it behaves as expected. For different size $L \times L$ toric codes we randomly apply X errors to the physical qubits with probability p. These errors are decoded and each time we check whether or not the decoder has succeeded. This is repeated for $N = 1000$ samples and the fraction that succeed gives the probability of successfully decoding P_{succ}. Figure 2.6 plots P_{succ} against the physical error rate p. We see evidence of a critical threshold error rate $p^* \approx 8.4\%{-}8.5\%$.

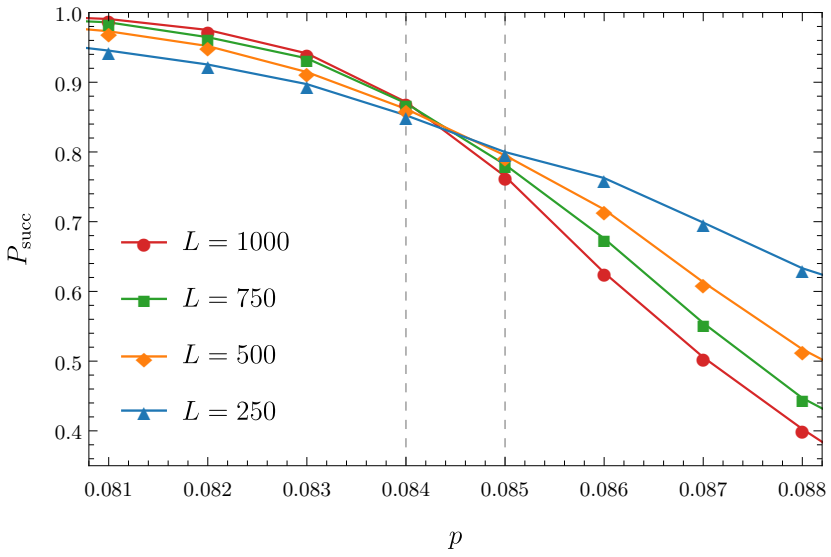


Figure 2.6. – Decoder success rate when attempting to correct i.i.d. noise as a function of the physical error rate p for different system sizes. We see the expected crossover behaviour for increasing system size and can identify a critical error rate approximately in the range 8.4% to 8.5%.

This is indicated by the crossing of the curves corresponding to different system sizes [78]. For physical error rates $p < p^*$ the probability of success P_{succ} increases with increasing system size, since larger sizes act as an increasingly strict constraint on which errors of a given density can cause the decoder to fail. Whereas for $p > p^*$ the opposite is true, it becomes increasingly harder to find an error pattern that does not cause a logical error, and the system size scaling is reversed. A more careful analysis of the system size scaling would be needed to more precisely identify the value of the critical error rate [77].

2.5. Summary

This chapter has given an introduction to the quantum memories framework and described the simulation methods employed to study the toric code in the next chapter. This began with a review of the stabiliser formalism for quantum error correcting codes. Followed by a discussion of the related idea of quantum memory Hamiltonians, which seek to use energetic suppression of errors to protect encoded information. The second half of the chapter contained an in depth account of the master equation methods used to simulate finite temperature environments and the decoding procedure we will use to correct errors.

In the following chapter we study the coherence time of the toric code quantum memory.

The memory is initialised in a pure state and allowed to evolve in contact with a thermal environment. We seek to dissect the coherence time, so we consider different timescales at different points in the chapter. However, our primary definition of the coherence time τ is the expectation value of the time the decoder first fails. This is estimated by applying a clustering decoder to the thermal errors.

3.

Understanding coherence times of the toric code

In this chapter we carry out a numerical experiment to quantify and explain the coherence time of the toric code as a finite temperature quantum memory. It is well established that the two-dimensional toric code is a poor thermal memory. Rigorous upper bounds have been placed on its coherence time, which are constants independent of system size [21, 23]. However, it is still informative to study the mechanics of its thermal decoherence. We will find that the toric code shares some important features with more complex and thermally robust models, so that studying the toric code can help us to more clearly understand their coherence times. The key similarity is that we find different coherence time scaling in two different regimes. One for small system sizes and another for large systems, with these limits defined with respect to a thermal length scale that depends on temperature. The small system limit is characterised by a coherence time that does depend on the size of the system. Though in the case of the toric code this scaling is bad as increasing the system size makes the coherence time shorter. In contrast, the coherence time for large systems is independent of system size and only depends on temperature. This is the same generic behaviour seen in models exhibiting partial self-correction [25], though in those cases the finite size scaling for small systems is positive rather than negative.

The Hamiltonian for the toric code memory is given by eq. (2.3). For simplicity we take $\Delta_p = \Delta_v = \Delta$, though in fact, we only need to consider one species of excitations (as noted in the definitions of the decoder we use). We study the plaquette excitations. Finite temperatures are modelled and simulated as described in section 2.3 and the thermally introduced errors are decoded with the clustering decoder detailed in section 2.4. We use the statistical mechanics of two-dimensional random walks to describe and model the motion of thermally excited toric code anyons.

The crucial weakness of the toric code is that once a pair of anyons has been created, which has energy cost 2Δ, the anyons can diffuse with no additional energy cost. A clear

physical picture of the theoretical upper bound on coherence time is that it is the time we expect to wait for the creation of a single anyon pair, whose winding can create a logical error $\tau \lesssim e^{2\Delta\beta}$ [21, 23]. This agrees with the Arrhenius law heuristic. Our numerical results, however, show that the coherence time in the large system limit does not saturate this bound and is in fact exponentially worse, scaling as $\tau \sim e^{\Delta\beta}$.

3.1. Thermal excitations and random walks

At any temperature β there is a critical system size below which we only expect a single pair of excitations to be present at any time in a finite size $L \times L$ system. We can deduce this thermal length scale L^* by considering the density of anyons in equilibrium. In thermal equilibrium we would expect the average density of excitations to scale like $\rho \sim e^{-\Delta\beta}$. From this it follows that the number of anyon *pairs* present is

$$\langle N \rangle \sim \frac{L^2\rho}{2} = \frac{L^2 e^{-\Delta\beta}}{2}. \tag{3.1}$$

This gives a critical size $L^* \sim e^{\Delta\beta/2}$. For systems smaller than this we have $\langle N \rangle \lesssim 1$, meaning the probability that there is more than a single pair of anyons on the lattice is negligible. The thermal decoherence of encoded information will then occur as the result of the behaviour of just this single pair of anyons. In contrast, for larger system sizes (or higher temperatures for the same system size) where $\langle N \rangle \gg 1$ many anyon pairs are present and we expect the interplay of the separate pairs to dominate the coherence time of the system. These different system size regimes of the toric code are illustrated in fig. 3.1. We will deduce estimates for the coherence time in both these limits and then compare these estimates to numerics. We first consider the small system limit before moving on to large systems.

3.1.1. Small system limit

For small systems we typically expect a single pair of excitations to cause the toric code memory to fail. We estimate the coherence time as $\tau_{\text{small}} = \tau_c + \tau_m$, where τ_c is the time it takes for this pair to be created. We denote by τ_m the time it takes for the anyons to diffuse

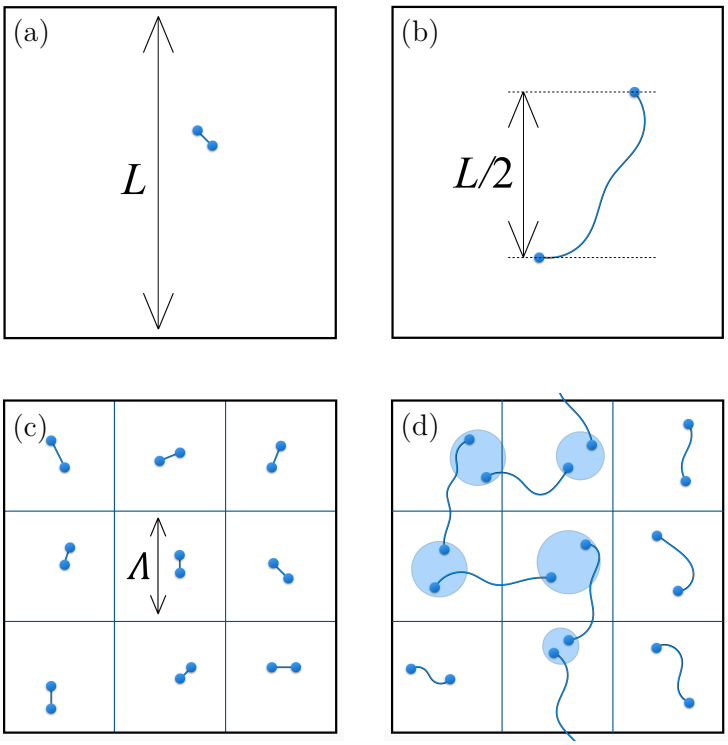


Figure 3.1. – Typical string-like error configurations for the different system size limits of the toric code at finite temperature. Configurations that cause the decoder to fail are also depicted. (a) when the system is small we only expect a single pair of anyons. (b) the decoder fails when the single pair is separated by more than $L/2$. (c) for large sizes the system contains many pairs with average spacing $\Lambda = \rho^{-1/2}$. (d) the decoder begins to mismatch pairs, failing when an error chain percolates.

across the lattice up to a separation $^{L}/_{2}$, such that encoded information is irrecoverable by quantum error correction. We show such an un-decodable configuration in fig. 3.1(b). Since we are only interested in the displacement between a single pair of anyons we treat one member of the pair as fixed with the other moving relative to it. When the anyons are not neighbours they move independently as an unbiased random walk. Using this description we can estimate the diffusion time τ_m.

The typical number of steps required for a two-dimensional random walker to reach a distance $^{L}/_{2}$ from its starting point is $(^{L}/_{2})^2$. Using the noise model described in section 2.3 the timescale of random walk steps is $^{1}/_{8\gamma(0)}$, where the factor of 8 counts the number of processes on the lattice that move anyons. We thus obtain

$$\tau_m \simeq \frac{1}{8\gamma(0)} \left(\frac{L}{2}\right)^2 = \frac{\beta L^2}{32\kappa}. \tag{3.2}$$

To estimate τ_c we note that the energy cost of pair creation is 2Δ. Again applying the noise model from section 2.3, pairs are created from the vacuum at rate $R_0 = 2L^2\,\gamma(-2\Delta)$. This implies the time we wait to see a creation event is $^{1}/_{R_0} \sim e^{2\Delta\beta}/(\kappa\Delta L^2)$. However, not all pairs diffuse to the required distance. Some creation pairs will instead fuse back to the vacuum (self-annihilate) at some point later in time. By analysing the random walks the pairs make we can quantify the effect of these self-annihilations on the dynamics.

We wish to estimate the probability that a pair does not self-annihilate before reaching separation $^{L}/_{2}$, we denote this as $\Pi(L,\beta)$. The motion of the pair (when the anyons do not neighbour each other) is described by an unbiased random walk. The probability that a two-dimensional random walk on a square lattice does not return to its origin in the first n steps scales as $^{1}/_{\ln(n)}$. On average we need the walker to avoid self annihilation for $(^{L}/_{2})^2$ steps. We can therefore expect a factor of $[2\ln(^{L}/_{2})]^{-1}$ in Π. Additionally, we know that when the anyons neighbour each other there is a temperature dependent bias towards self-annihilation since $^{\gamma(2\Delta)}/_{\gamma(0)} \sim \Delta\beta$. The probability that they separate rather than annihilate should then include a factor $\sim 1/(1+A\Delta\beta)$ to account for this bias, where A is some constant and $A\Delta\beta$ is the relative chance the anyons annihilate when they are

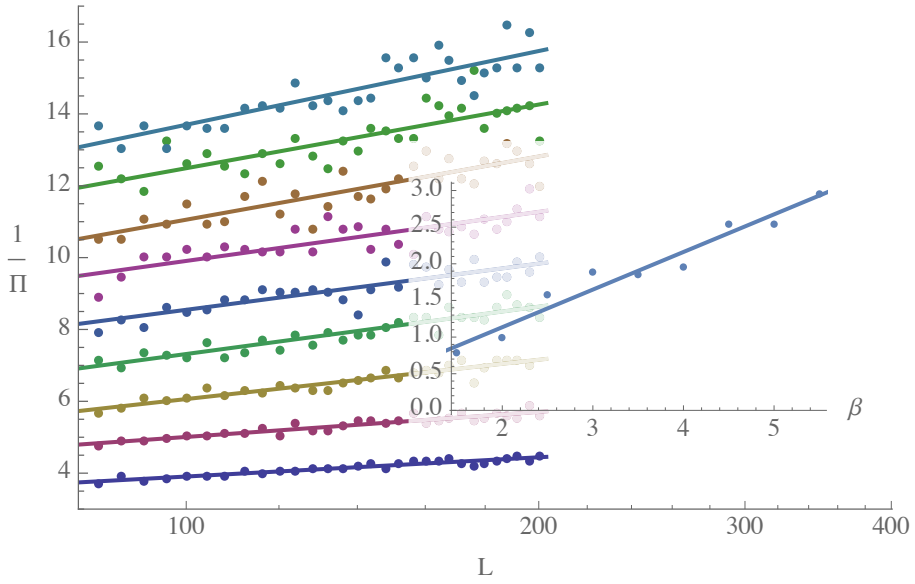


Figure 3.2. – Log-linear plot of the reciprocal of Π plotted against L, for a range of β values. We find $^1/_\Pi$ grows linearly with $\ln(^L/_2)$. The Inset shows the gradient of the fits in the main plot as a function of β.

nearest neighbours. The probability $\Pi(L, \beta)$ then takes the form

$$\Pi(L, \beta) \sim \frac{1}{1 + A\Delta\beta} \frac{1}{\ln\left(^L/_2\right)}. \tag{3.3}$$

Combining these elements, we expect a creation timescale

$$\tau_c \simeq \frac{1}{R_0} \frac{1}{\Pi} \sim \frac{e^{2\Delta\beta}}{\kappa\Delta L^2} (1 + A\Delta\beta) \ln\left(^L/_2\right). \tag{3.4}$$

The total coherence time is $\tau_{\text{small}} = \tau_c + \tau_m$. In the small-system limit time τ_c is the dominant contribution to τ_{small} due to its exponential dependence on β.

To test these predictions we study the system evolving in this limit using different numerical experiments with various initial conditions and employing some variations to the physical noise model. We separately estimate $\Pi(L, \beta)$, τ_c and τ_m using numerical simulations. To compute the probability $\Pi(L, \beta)$ we alter the standard simulation scheme; rather than beginning in a ground state we initialize the system with a single pair of anyons present on the lattice. These initialised anyons are nearest neighbours on the lattice mirroring the situation after a thermal creation event. Additionally, we set the rate of creation equal to zero such that no additional pairs of anyons are created. In this way

we model only the random motion of a single pair of anyons walking across the lattice. We evolve the system until either the separation of the anyons reaches the Euclidean distance $L/2$, or the pair meets at a common point on the lattice and annihilates. The quantity Π is estimated as the fraction of samples that reach separation $L/2$ rather than annihilating, we estimate by sampling over 10^4 trials. The results we obtain are shown in fig. 3.2. The fit we find to the data is

$$\frac{1}{\Pi} = (0.108 + 0.513\Delta\beta)\ln(L/2). \tag{3.5}$$

This is in good agreement with the scaling form we hypothesised in eq. (3.3), where constant $A \approx 5$.

We next numerically investigate the scaling of τ_c and τ_m with L and β. Timescale τ_c is the typical wait time for the creation of the pair of anyons that will cause a logical error and τ_m is the average length of time from their creation until they reach a logical distance, $L/2$, from each other. To find these values we simulate the thermal evolution of the system prepared in a ground state. We attempt to decode the lattice after each event that occurs due to the noise simulation, and the simulation is stopped once the decoder fails. Every time a pair of anyons is created its creation time is recorded, this information is discarded if the pair later annihilate. We test the prediction that the environment never creates more than a single pair of anyons at a given time and see it to be valid. Estimates of the creation timescale τ_c and diffusion time τ_m are obtained by averaging over 10^3 simulations. The data for each is plotted in fig. 3.3 and fig. 3.4 respectively. Our data correlates well with the forms we expect, which are eq. (3.2) and eq. (3.4). We find numerical fits

$$\kappa\Delta\tau_m \simeq 0.028\Delta\beta L^2 \quad \text{and} \quad \kappa\Delta\tau_c = 0.150\frac{e^{1.99\Delta\beta}}{L^{2.01}}\frac{1}{\Pi}\,,$$

where we take the function Π to be the one we evaluated numerically, eq. (3.5).

The Arrhenius law estimate of the toric code coherence time $\tau \sim e^{2\Delta\beta}$. We have shown that in the small size regime the leading contribution to the coherence time is τ_c, given by eq. (3.4). If we neglect the sub exponential β dependence inside $1/\Pi$, we can approximate

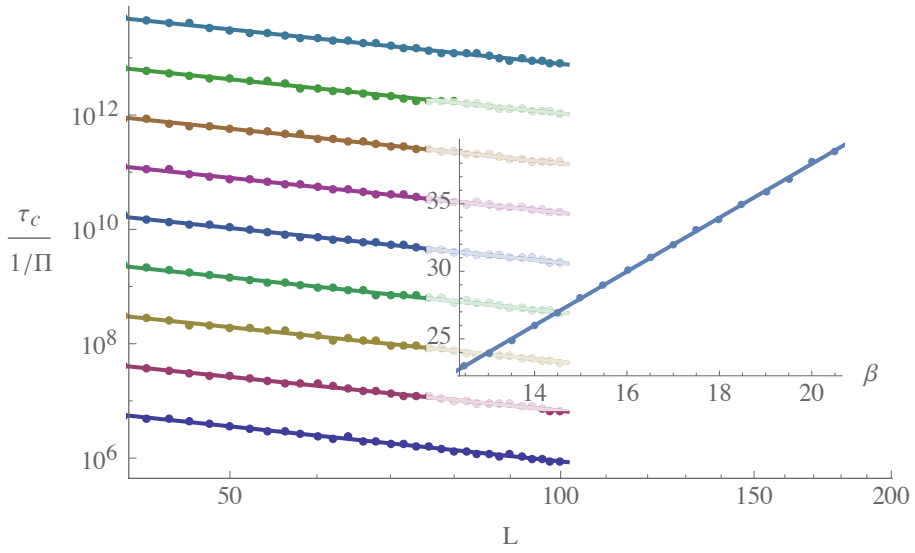


Figure 3.3. – Log-log plot of the contribution to the coherence time from pair creation τ_c as a function of L for a range of β. We expect τ_c to include a factor of $1/\Pi(L,\beta)$, which we have separately found numerically, so we divide this away. Estimates are obtained over 1000 simulations. The fits are linear and all have similar gradient with mean of -2.01. Inset is the intercept of the linear fits as a function of β. Overall we see a dependence $\tau_c = 0.150(e^{1.99\beta}/L^{2.01})(1/\Pi)$, which is in excellent agreement with the predicted behaviour given in eq. (3.4). In our data Δ and κ are both separately set equal to one. Equivalently, τ_c is measured in units of $1/(\kappa\Delta)$ and β in units of $1/\Delta$.

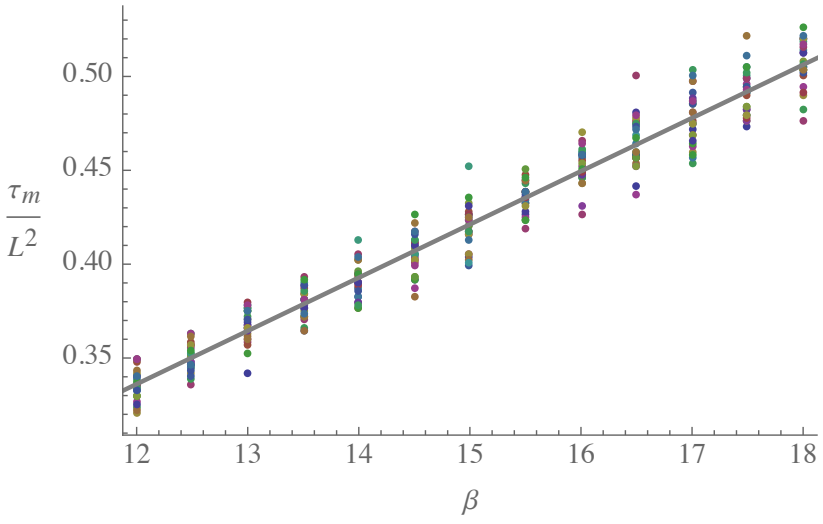


Figure 3.4. – Contribution to the coherence time from anyon motion τ_m plotted against β for a range of system sizes L =50, 52, 54, ..., 100. By dividing τ_m by L^2 we see data collapse of different system sizes. The linear fit displayed uses the average values of the fit parameters obtained from the separate system sizes $\tau_m \simeq 0.028\,\beta L^2$. Again Δ and κ are both separately set equal to one, which is equivalent to measuring τ_c in units of $1/(\kappa\Delta)$ and β in units of $1/\Delta$.

the coherence time by

$$\tau_{\text{small}} \sim e^{2\Delta\beta} \frac{\ln(L/2)}{L^2}. \tag{3.6}$$

We see that as L becomes larger the coherence time of the toric code decreases polynomially with L up until the critical size scale $L^* \sim e^{\Delta\beta/2}$. At this point the system transitions from the small size limit into the large size regime. We next show that for large systems the lifetime loses any dependence on system size and is determined only by β.

3.1.2. Large system limit

In the large size limit thermalisation creates many anyons. In this regime the equilibrium density of single anyons is $\rho \sim e^{-\Delta\beta}$. On average, anyon pairs are created uniformly throughout the system and each can expect to solely occupy an area of approximately $2/\rho$. We approximate this area as a square of linear size $\Lambda = (\rho/2)^{-1/2} \sim e^{\Delta\beta/2}$ as shown in fig. 3.1(c). The probability the decoder fails becomes appreciable once some fraction of pairs separate to distance Λ such that an error chain can percolate, as illustrated in fig. 3.1(d).

Within a single $\Lambda \times \Lambda$ region the evolution proceeds like the small size case, i.e. there will be a creation event and the anyons subsequently diffuse apart. We consider a region as having failed once its anyons move closer to anyons from neighbouring regions than to each other. Assuming each region evolves independently, the time the system decoheres τ_{large} is estimated by the time the typical region fails. This is given by eq. (3.6) where we set $L/2 = \Lambda$. Applying eq. (3.2) and eq. (3.4) where we also set $L/2 = \Lambda$ gives the diffusion and creation timescales $\tau_{m,\Lambda} \sim \beta\Lambda^2/\kappa \simeq \beta e^{\Delta\beta}/\kappa$ and

$$\tau_{c,\Lambda} \sim \frac{e^{2\Delta\beta}}{\kappa\Delta\Lambda^2}(1 + 5\Delta\beta)\ln(\Lambda). \tag{3.7}$$

To get an expression in terms of β we write $\Lambda = C\, e^{\Delta\beta/2}$, where C accounts for the constants absorbed into ρ. Then substituting Λ into eq. (3.7) and rearranging

$$\tau_{c,\Lambda} \sim \frac{e^{\Delta\beta}}{\kappa\Delta}(1 + r\Delta\beta + s\Delta^2\beta^2), \tag{3.8}$$

where r and s are new constants related to those already introduced. In contrast to

the small size case the timescales for creation and diffusion have the same exponential dependence on β, therefore we expect them both to contribute appreciably to the coherence time. Combining terms we obtain high-temperature coherence time

$$\tau_{\text{large}} \simeq \tau_{c,\Lambda} + \tau_{m,\Lambda} \simeq \frac{e^{\beta\Delta}}{\kappa\Delta}(1 + r'\Delta\beta + s\Delta^2\beta^2)\,, \tag{3.9}$$

where r' also includes the motion contribution. This simple modelling ignores thermal effects that can now take place such as pair fusion, where two anyons from different pairs fuse to vacuum combining their strings to form a longer string. Processes such as these will alter both the diffusion speed and the pair self-annihilation probability. As a result we do not expect this model to give good predictions of the values of r' and s. However, we do expect it to correctly predict the general features of the high temperature dynamics. In particular the altered exponential dependence on β and the system size independence of the lifetime. If we ignore its sub-exponential β dependence the large size coherence time is approximated by

$$\tau_{\text{high}} \sim e^{\Delta\beta}\,. \tag{3.10}$$

This is also exponentially growing with β but at a much slower rate than the Arrhenius law would predict.

We simulate the system evolving in this regime in order to test our assumptions and verify our predictions about the dynamics. First, we confirm that the anyon densities at the time the decoder fails obeys $\rho \sim e^{-\Delta\beta}$ and that the typical number of anyons is always large, $\langle N\rangle \gg 1$, for the parameter ranges we consider. In order to test that the average motion is the main decoherence mechanism, rather than rare anyon pairs that move fast and separate to logical distances $L/2$, we probe the distance between anyon pairs at the time the decoder fails. During the simulation anyons that are created together are marked as a member of their given creation pair. At the end of the simulation we measure the Euclidean distance between each marked pair. If two anyons from two separate pairs fuse, the remaining two unpaired anyons on the lattice are marked as members of the same pair. As one might expect, we find that pairs created from a fusion of two separate pairs do typically achieve a greater separation. However, this effect is small with respect to the average data. Our numerical results show that the average separation between all anyon

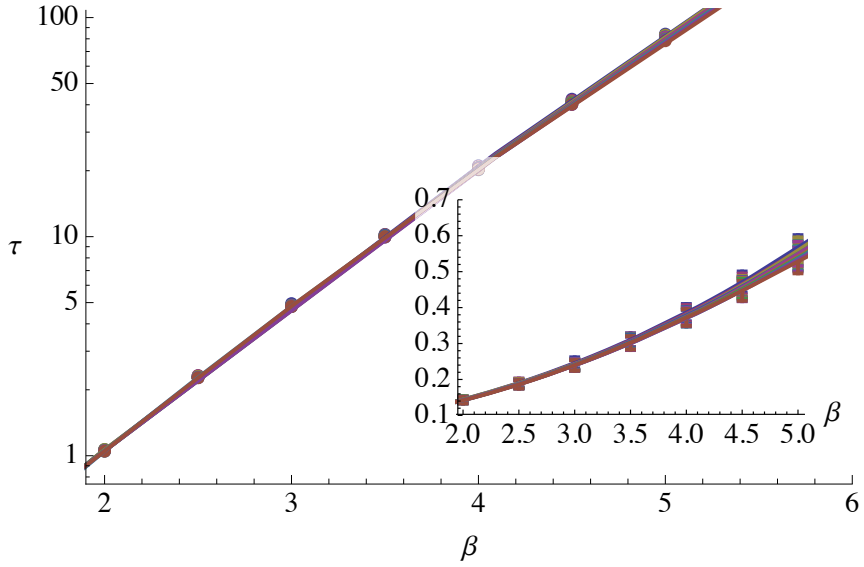


Figure 3.5. – Coherence time of the large size toric code as a function of β for a range of system sizes L =100, 120, ... 200. Estimates are averages over 1000 simulations. Shown are fits to the form of eq. (3.9). Inset is the non-exponential part of the scaling obtained by dividing off e^{β}. We observe only small variations in the fit parameters between system sizes and their averages give: $\tau_{\text{high}} = 0.56\, e^{1.01\beta}(1 + 0.28\beta + 0.31\beta^2)$. Here Δ and κ have both been separately set equal to one, which is equivalent to measuring τ in units of $1/(\kappa\Delta)$ and β in units of $1/\Delta$.

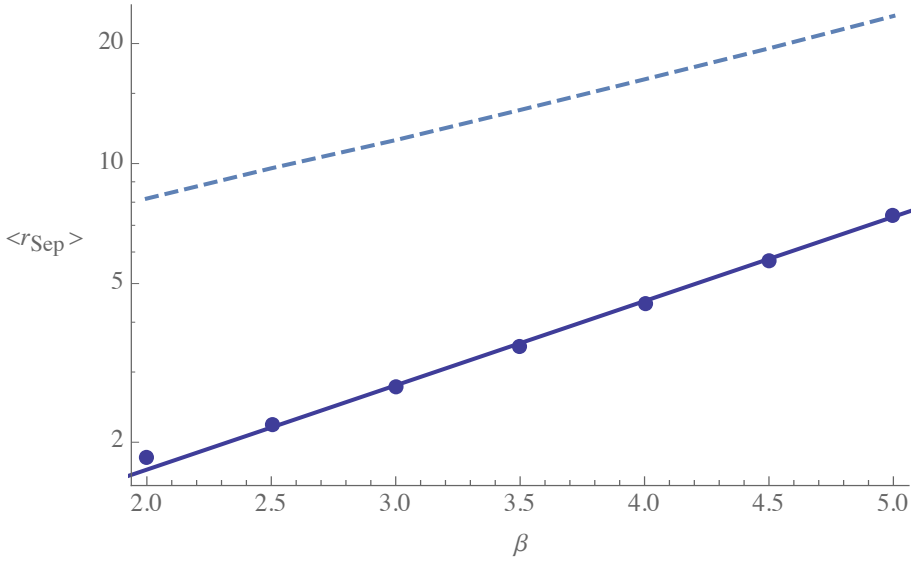


Figure 3.6. – Log plot of average separation between paired anyons at the time the decoder fails as a function of β, for system size $L = 200$. This is typical of our results for different system sizes. The solid line is a fit to the data with linear gradient approximately 0.49 in this Log plot, consistent with scaling $\langle r_{\text{sep}} \rangle \sim e^{\Delta\beta/2}$. The dashed line is the maximum value of the separation, averaged over the simulation runs, which is seen to be much less than $L/2$. In this plot Δ is set equal to one, which is equivalent to measuring β in units of $1/\Delta$.

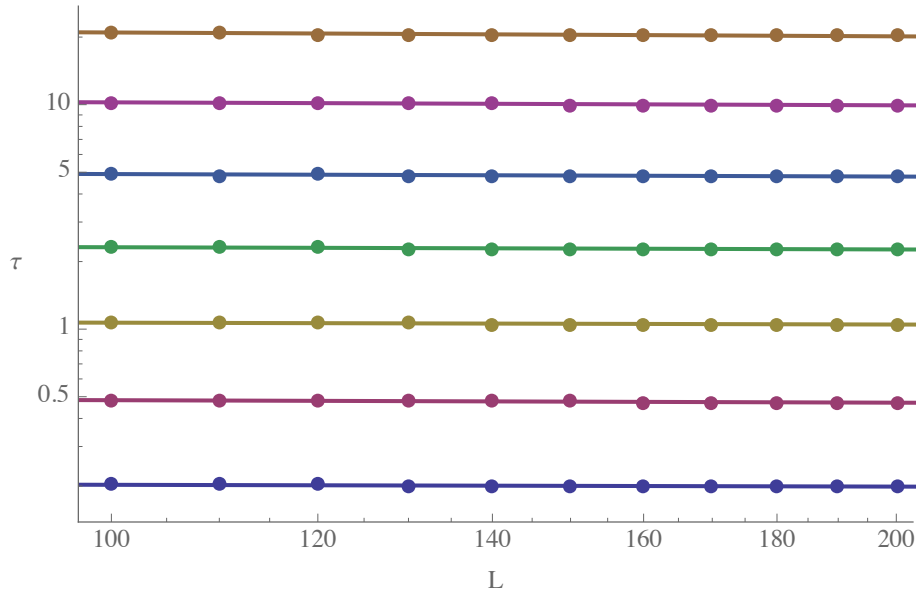


Figure 3.7. – Approximate system size independence of the large system size coherence time. Plots of τ on a Log scale against L for different temperatures $\beta = 2.0, 2.5, \ldots 5.0$. Linear fits show very small negative gradients ranging from -0.03 for the lowest value of β, down to -0.06 for the highest β. However, zero gradient is well within the 95% confidence interval of each fit. Again Δ and κ are both separately set equal to one, or equivalently τ is measured in units of $1/(\kappa\Delta)$.

pairs grows like $e^{\Delta\beta/2}$, as predicted. In addition we find that the maximum separation between any pair is always much less than $L/2$. The scaling of the average and maximum separations with β are shown in fig. 3.6. These observations confirm that the decoherence results from the average motion of anyons in local regions.

As with the small system case, to find the coherence time we begin with the system in the ground state and evolve it under the thermal noise model. We attempt to decode after each time step and the coherence time is computed as the mean time the decoder first fails. Figure 3.5 shows the coherence time as a function of β. We see a clear dependence on $e^{\Delta\beta}$ with a sub-exponential quadratic component as predicted by eq. (3.9). In fig. 3.7 the coherence time is plotted as a function of system size and we see strong evidence that τ is independent of L.

3.2. Conclusions

We have examined the thermal motion of toric code excitations that lead to a failure of the toric code as a quantum memory. The dynamics were modelled using the statistical

mechanics of two-dimensional random walks. We identified two regimes with distinct scaling behaviours of the coherence time. Small systems are characterised by a coherence time that depends on temperature β and system size L through

$$\tau \sim e^{2\Delta\beta} \frac{\ln(L/2)}{L^2} \,. \tag{3.11}$$

This is the behaviour expected from the Arrhenius law, corrected by a finite size factor. The finite size correction encodes the entropic drivers of decoherence. Namely, as the system grows larger there are more physical qubits that can suffer the initial pair creation error. Thus causing the coherence time to decrease with volume as $1/L^2$. This is offset by the increased difficulty now faced by the anyon pair of completing a random walk that causes a failure. However, this defensive effect has much weaker dependence on system size only growing with $\ln(L)$. The net result is that growing entropy causes the coherence time to decrease with increasing system size. Above a critical system size $L^* \sim e^{\Delta\beta/2}$ the scaling of τ transitions towards its large size behaviour. In this limit τ becomes independent of L and scales as $\tau \sim e^{\Delta\beta}$. At the critical L^* the entropy becomes strong enough to create many anyonic excitations in the system. These diffuse, overlap and the error percolates causing decoherence. The same entropic considerations dominate the failure of partially self-correcting memories above a critical size. The behaviour predicted by our analysis was tested with numerical simulations and we found it to be in good agreement.

4.

Background 2:
Kitaev's Honeycomb

This chapter introduces the Kitaev Honeycomb model and outlines its solution, as originally obtained by Kitaev [51]. In section 4.1 the model is defined in terms of a lattice of spins. Section 4.2 then describes Kitaev's fermionisation scheme that represents the spins as Majorana fermions. This leads to an efficiently solvable problem of superconducting fermions on a lattice, giving an explicit form for the states of the model. A complete account of diagonalising the fermionic degrees of freedom is given in appendix A.

4.1. Spin Model

The model is defined for spins on a honeycomb lattice. Spin-1/2 particles are placed at the vertices, and interact anisotropically with their nearest neighbours. Weak three-body interactions are also included in the Hamiltonian to explicitly break time-reversal symmetry. These three-body interactions can be justified as the perturbation coming from a weak magnetic field [51] or, in the context of quantum simulators, as an engineered interaction [80, 81]. The interaction terms do not commute with each other making the model highly frustrated and in general very difficult to solve. However, it has an extensive number of conserved quantities associated with stabiliser-like operators on all the hexagonal lattice plaquettes that can be exploited to find a solution.

To describe the Hamiltonian the links of the honeycomb lattice are divided into three types ($\alpha = x, y, z$) based on their spatial orientation. Each of these links is associated with a two-body interaction term that acts on the two spins at the end of the link with Pauli operators σ_α and strength J_α, as illustrated in fig. 4.1(a). Note that, following the convention used in literature on the Kitaev Honeycomb, we have changed notation from the previous chapters and the symbols σ_x, σ_y and σ_z will now be used to denote the Pauli

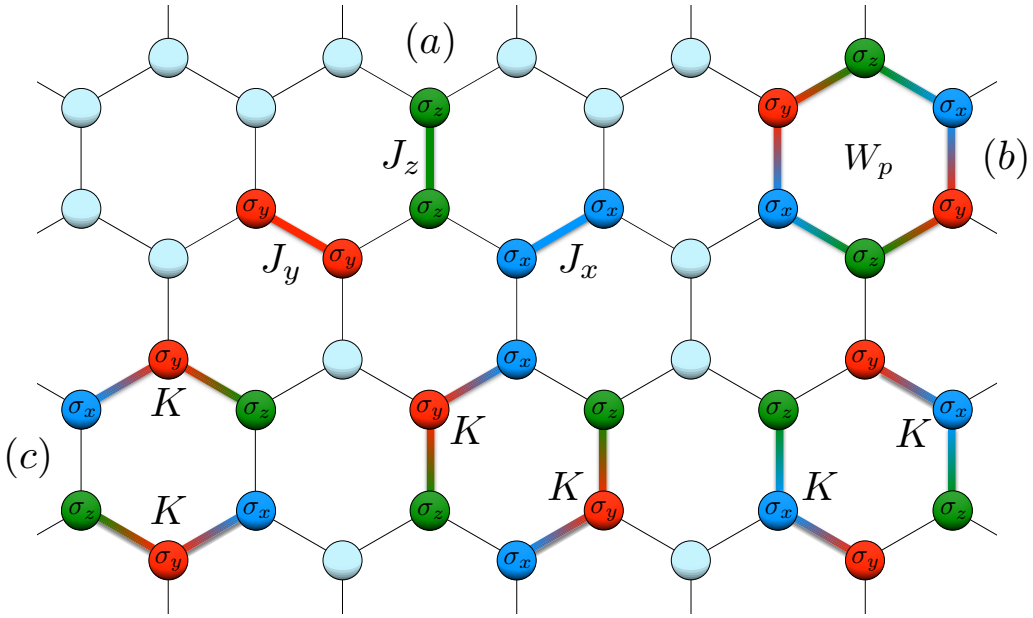


Figure 4.1. – The Kitaev Honeycomb in the spin-picture. Spin-$^1/_2$ particles are located at the vertices of a Honeycomb lattice. (a) two-body terms of the Hamiltonian with strengths J_x, J_y and J_z act along x-type, y-type and z-type bonds. (b) the six-body plaquette operator W_p that commutes with all terms in the Hamiltonian. (c) three-body Hamiltonian terms with strength K.

matrices, which are given explicitly by

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Similar to the notation used previously, σ_x^i means an operator that acts on spin i with a Pauli σ_x and on all the other spins with the identity. The product of the two-body operators around each plaquette p generates a Hermitian operator

$$\begin{aligned} W_p &= (\sigma_y^{12\text{ o'clock}}\sigma_y^{2\text{ o'clock}})(\sigma_z^{2\text{ o'clock}}\sigma_z^{4\text{ o'clock}})(\sigma_x^{4\text{ o'clock}}\sigma_x^{6\text{ o'clock}})\ldots \\ &\quad\times(\sigma_y^{6\text{ o'clock}}\sigma_y^{8\text{ o'clock}})(\sigma_z^{8\text{ o'clock}}\sigma_z^{10\text{ o'clock}})(\sigma_x^{10\text{ o'clock}}\sigma_x^{12\text{ o'clock}}) \\ &= \sigma_z^{12\text{ o'clock}}\sigma_x^{2\text{ o'clock}}\sigma_y^{4\text{ o'clock}}\sigma_z^{6\text{ o'clock}}\sigma_x^{8\text{ o'clock}}\sigma_y^{8\text{ o'clock}} \end{aligned} \tag{4.1}$$

where we have used the reference of points on a clock face to index the spins around the plaquette. This plaquette operator is illustrated for a single plaquette in fig. 4.1(b). Three-body interaction operators are chosen so as to commute with these W_p. This allows six different types of three-body operators, each illustrated in fig. 4.1(c). These are copied

and pasted everywhere on the lattice. The Hamiltonian can then be written

$$H = -J_x \sum_{(i,j):x} \sigma_x^i \sigma_x^j - J_y \sum_{(i,j):y} \sigma_y^i \sigma_y^j - J_z \sum_{(i,j):z} \sigma_z^i \sigma_z^j - K \sum_{(i,j,k)} \sigma_x^i \sigma_y^j \sigma_z^k \,, \tag{4.2}$$

where the notation $(i,j):\alpha$ means the sum runs over all α type links with i, j being the spins at the ends of the link and the triples of the three-body terms (i,j,k) are described previously. The coupling constants J_α set the strength of each of the nearest-neighbour interactions and $K \ll J$ sets the strength of the three-body terms.

The plaquette operators W_p commute with the Hamiltonian as well as with each other,

$$[H, W_p] = 0 \text{ for all } p$$
$$[W_p, W_q] = 0 \text{ for } p \neq q \,,$$

meaning the eigenvalues of the W_p are constants of motion. Therefore, H is block diagonal in the basis of the W_p. A clever transformation of the spins to fermions explicitly separates out these plaquette degrees of freedom and solves the model (up to diagonalising a matrix that is polynomially big in the number of spins). There are different ways to carry out the fermionisation [51, 82–84] but here we describe Kitaev's approach [51].

4.2. Exact solution

This section describes Kitaev's solution to the model. This involves a transformation from spins to Majorana fermions, so the section starts by defining these fermions. Following this, the steps of the transformation are given as well as the translation of all important quantities from the spin picture to the Majorana representation. This ends with a discussion of the interpretation of the Majorana representation in terms of vortices and gauge fields. The diagonalisation of the fermionic degrees of freedom is then shown. The section ends with a short review of the properties of the solutions.

4.2.1. Majorana fermions

The transformation represents the original spins with a particular type of fermion called Majorana fermions. These are *real* fermions described by operators c with the properties:

- hermitian $c = c^\dagger$,
- square to the identity $c^2 = c^\dagger c = \mathbb{1}$,
- pairwise anti-commute $c\,c' = -c'c$.

These Majorana are distinct from *complex* fermions (such as those we would conventionally use to describe electrons in solids), which have separate creation and annihilation operators $a^\dagger$ and a. But the two representations are interchangeable, any complex fermion can be described by two Majorana c and c' through the transformation

$$a = \frac{c + ic'}{2}, \qquad a^\dagger = \frac{c - ic'}{2}\,. \tag{4.3}$$

4.2.2. Kitaev's fermionisation

Under Kitaev's fermionisation each spin site, fig. 4.2(a), is represented by four Majorana operators c_i, b_i^x, b_i^y and b_i^z, fig. 4.2(b). The Hilbert space acted on by these four operators is redundantly large for representing the spins. Two Majorana describe a single complex fermion whose Hilbert space has the same dimensions as a spin-$^1/_2$ particle, so the Hilbert space of the four operators $\{c_i, b_i^x, b_i^y, b_i^z\}$ is twice as large as is needed. A solution to this problem is to specify half the states spanned by $\{c_i, b_i^x, b_i^y, b_i^z\}$ as *physical* and the rest as unphysical. An operator that divides the Majorana Hilbert space in half is the total fermionic parity

$$D_i = b_i^x\, b_i^y\, b_i^z\, c_i\,. \tag{4.4}$$

Arbitrarily, the eigenstates of D_i with $+1$ eigenvalue are chosen as the physical states. The Pauli operators of the spins are expressed in terms of the Majorana as

$$\tilde{\sigma}_\alpha = i b^\alpha c\,,$$

where these $\tilde{\sigma}$ obey the correct algebra when restricted to the physical subspace.

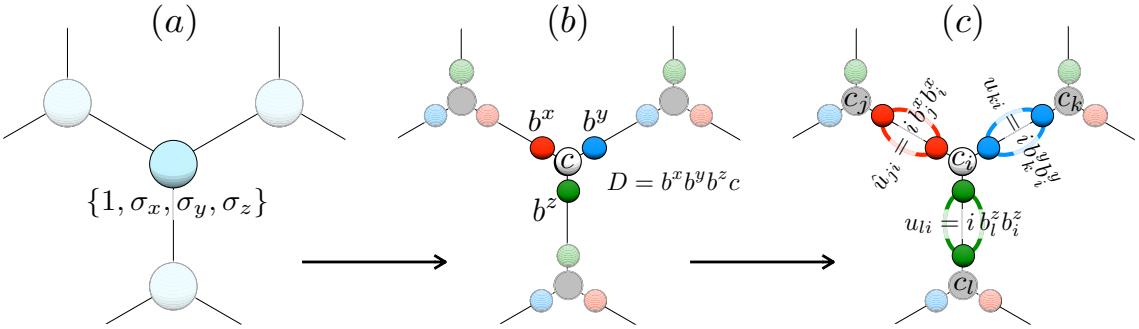


Figure 4.2. – Transformation from spins to Majorana. (a) the original spins of the Honeycomb model, described by Pauli operators $\{1, \sigma_x, \sigma_y, \sigma_z\}$. (b) each spin is represented by four Majorana modes b^x, b^y, b^z and c. Physical states in this Majorana representation are eigenstates of operator D with eigenvalue $+1$. (c) link operators $\hat{u}_{ji}$ measure the fermionic parity of pairs of $x/y/z$ Majorana along $x/y/z$ type links between nearest-neighbour spin sites.

The c_i Majorana become the physical degrees of freedom living at each lattice site. In contrast, the b_i^α Majorana are hybridised along links to form fermionic parity observables. These link operators are defined by $u_{ij} = i b_i^\alpha b_j^\alpha$ for each of the x, y and z links, as illustrated in fig. 4.2(c). The u_{ij} are anti-symmetric, $u_{ji} = -u_{ij}$, and Hermitian. Additionally, they square to the identity meaning they have eigenvalues ± 1. An efficient way to encode a pattern of u_{ij} eigenvalues is to draw the lattice links as arrows as in fig. 4.3, such that the arrowhead indicates the ordering (i, j) that gives $+1$ eigenvalue. In the Majorana representation the plaquette observables W_p are expressed as products of these link operators. Let p be the eigenvalue of W_p and l_{ij} the eigenvalue of u_{ij} then

$$p \;=\; \prod_{(i,j)\,\in\,\partial p} l_{ij}\,. \tag{4.5}$$

The sequence $(i, j) \to (j, k) \to \ldots$ should hop in one direction around the plaquette to avoid any possible ambiguities from the anti-symmetry of the the u_{ij}.

The Majorana Hamiltonian $\tilde{H}$ is written in terms of the lattice Majorana c_i and the link operators u_{ij} and has the form

$$\tilde{H} \;=\; \frac{i}{4}\sum_{i,\,j}\left(2J_{ij}\,u_{ij}\,c_i\,c_j + 2K\sum_k u_{ik}\,u_{jk}\,D_k\,c_i\,c_j\right) \;\equiv\; \frac{i}{4}\sum_{i,\,j} A_{ij} c_i c_j\,. \tag{4.6}$$

In the physical subspace $\tilde{H}$ is equivalent to the spin Hamiltonian, eq. (4.2). This Hamil-

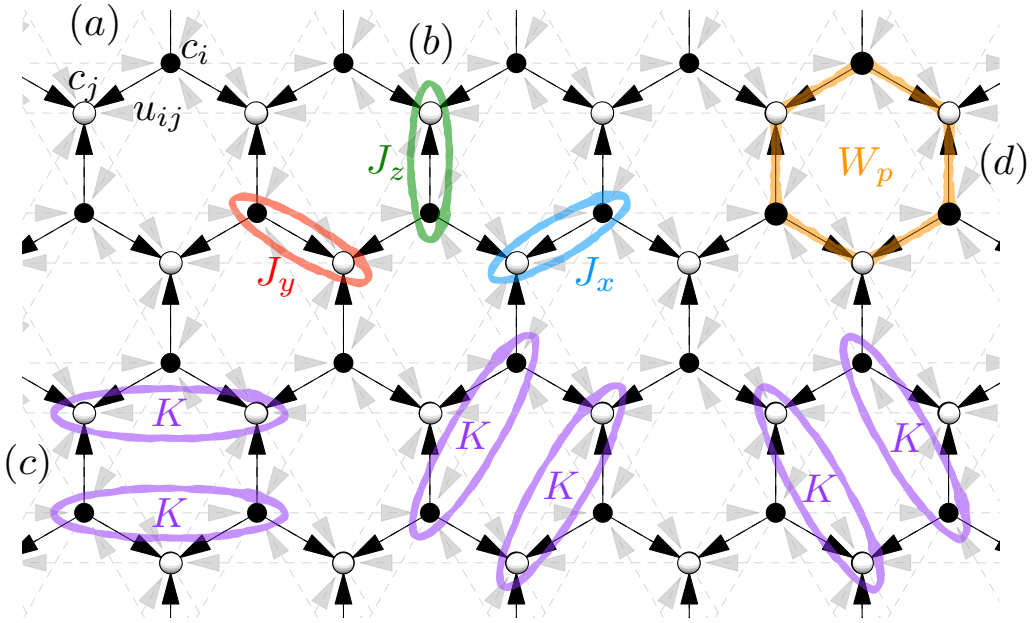


Figure 4.3. – Majorana representation of the model in the no-vortex sector, based on Kitaev's fermionisation. (a) lattice Majorana c and link operators u_{ij}. The arrow indicates the (i, j) ordering that gives $+1$ eigenvalue. (b) two-body spin interactions become nearest-neighbour Majorana parity weights. (c) three-body spin interactions give next-nearest-neighbour Majorana parity weights. (d) the plaquette operators become the Wilson loops of the u_{ij} gauge field.

tonian is still hard to solve, however, as the c_i are coupled to the dynamic u_{ij} degrees of freedom making the model still highly interacting. The next step is to use the physical symmetries to restrict to particular subspaces where the u_{ij} become static and the lattice of Majorana becomes free and easy to solve.

4.2.3. Vortex sectors and gauge fixing

As discussed in the spin picture, the plaquette operators W_p commute with the Hamiltonian and hence $\tilde{H}$ is block diagonal in the basis of W_p eigenstates. These blocks, labelled by the eigenvalues of the W_p operators, are called the *sectors* of the model. In addition to the plaquette observables, $\tilde{H}$ commutes with the u_{ij} so the link operators are local symmetries within the Hilbert space of the $\{c_i, b_i^x, b_i^y, b_i^z\}$ Majorana. However, comparing this back to the spin picture these conserved link variables do not seem to correspond to any spin constants of motion. This complication has arisen from the fact the u_{ij} do not commute with the physical subspace operators D_i. The u_{ij} anti-commute with the D_i meaning the physical states do not have well defined values of u_{ij}. In fact, it can be

shown that the spin states obtained by projecting from a state of well defined u_{ij} into the physical subspace,

$$|\Psi_{\text{spin}}\rangle = \mathcal{P} \left|\Psi_{u_{ij}}\right\rangle \quad \text{where} \quad \mathcal{P} = \prod_i \frac{\mathbb{1} + D_i}{2}\,, \tag{4.7}$$

become superpositions of many u_{ij} patterns [85]. However, the physical observables such as energy, $\tilde{H}$, and W_p do commute with D_i and so are invariant under applying projector $\mathcal{P}$. This observation gives an interpretation of the u_{ij} as a $\mathbb{Z}_2$ lattice gauge field, with the W_p the associated observable Wilson loops. Negative, -1, eigenvalues of W_p signify that a π flux, or vortex, lies on the plaquette. Negative eigenvalues of u_{ij} can be interpreted as an unphysical string passing perpendicularly through the edges of the lattice. These open strings must terminate on π fluxes. The D_i fulfil the role of local gauge transforms of the u_{ij} field. Specifically they act to locally deform these strings around vertices. Particularly important sectors of the model are often named with reference to this perspective. The uniform $W_p = +1$ configuration is called the *no-vortex sector*. By Lieb's theorem the no-vortex sector contains the ground state [51, 86]. We will also reference the uniform $W_p = -1$ configuration, which is called the *full-vortex sector* [53]. These are amongst many others that have been identified as having interesting properties [54].

The realisation that the u_{ij} act as a gauge field to the W_p, with classical configurations of the two linked through eq. (4.5), means that we can choose to restrict to a single sector by fixing the values of the u_{ij}. This choice will not be unique, but as long as the observables we want to evaluate using the fermionic states are gauge invariant (commute with D_i) the exact choice will not matter. Fixing u_{ij} values then transforms the Majorana Hamiltonian eq. (4.6) into a solvable problem of superconducting fermions on a lattice described by

$$\overline{H} = \frac{i}{4} \sum_{j,k=1}^{2L^2} A_{jk}\, c_j c_k\,, \tag{4.8}$$

where A_{jk} is now just a matrix of numbers rather than Hilbert space operators. This describes superconducting fermions since if we were to change into a basis of local complex fermions (for example by pairing the two Majorana along a z link) the Hamiltonian would contain both hopping and pairing terms of these complex fermions.

There are still some details of the gauge fixing and the projection procedure, eq. (4.7), that we have not discussed. These details will become important when considering the finite temperature states of the Honeycomb model, and we will return to them chapter 6. In the rest of this chapter we will outline the fermionic solutions of the superconducting Hamiltonian $\overline{H}$ and discuss some of its physical properties.

4.2.4. Normal modes

Following the gauge fixing described in the previous section, a complete solution for a finite size lattice of $L \times L$ plaquettes, is obtained by diagonalising the $2L^2 \times 2L^2$ Hermitian matrix (iA_{jk}) of eq. (4.8). The matrix A_{jk} is anti-symmetric and real valued. There are two important normal forms of A_{jk}, a diagonalisation using complex matrices and a block-diagonal form using real matrices, given respectively by

$$A = U\, D\, U^\dagger \tag{4.9}$$

$$A = Q\, \Sigma\, Q^T \tag{4.10}$$

where U is unitary, D diagonal, Q orthogonal and Σ is (2×2) block-diagonal. For an extended discussion of these matrices, and the relationships between them, see appendix A. (Numerically we prefer to solve for Q & Σ over U & D, see note B.1.1.) The eigenvectors of (iA_{jk}), the columns of U, encode the transformation from Majorana c_j to fermionic diagonal modes f_p, $f_p^\dagger$,

$$c_i = \sum_{p=1}^{L^2} \sqrt{2} \left([\underline{u}_p]_i\, f_p^\dagger + [\underline{u}_p]_i^*\, f_p \right) \tag{4.11}$$

where $\underline{u}_p$ is the eigenvector corresponding to the eigenvalue ε_p and the $(\ldots)^*$ in the second term is complex conjugation. Note that this sum has L^2 terms. The matrix (iA_{jk}) only has ${}^N\!/_2$ unique eigenvalues, which occur as plus/minus pairs. We choose $\varepsilon_p \geq 0$ to be the set of non-negative eigenvalues. An alternative representation of the energy eigenmodes is in terms of Majorana normal modes b'_p, b''_p given by

$$b'_p = f_p^\dagger + f_p\,, \quad b''_p = \frac{1}{i}(f_p^\dagger - f_p) \quad \text{where} \quad p = 1, \ldots, L^2\,.$$

(Analogous to eq. (4.11), the columns of matrix Q relate these Majorana normal modes to the lattice Majorana c_i.) Expanding the Hamiltonian in the energy eigenbasis, and alternatively in terms of the Majorana normal modes, it becomes

$$\overline{H} \;=\; \sum_{p=1}^{N/2} \varepsilon_p \left(f_p^\dagger f_p - \tfrac{1}{2} \right) \;=\; \frac{i}{2} \sum_{k=1}^{N/2} \varepsilon_p \, b'_p b''_p \,. \tag{4.12}$$

The energy eigenvalues of the fermions, ε_p, depend on the matrix A and on the vortex sector in particular. As previously mentioned, the ground state of the model corresponds to the no-vortex sector where all W_p have eigenvalue $+1$. For the full model H, there are two types of excitations above the ground state, fermions and vortices. Creating a quasi-particle excitation of $\overline{H}$ inserts a fermion into the system at an energy cost. Similarly, flipping the eigenvalues of a pair of the W_p by inserting an open gauge string creates two vortices. This new vortex sector has its own fermionic Hamiltonian $\overline{H}'$ with higher ground state energy than the no-vortex sector, describing the energy cost of creating vortices.

An important set of quantities characterising states of the fermions are the two-point Majorana correlation functions. In fact, simple states such as the ground state and the thermal state are completely characterised by these quantities. In the energy eigenbasis the two-point correlations are given by

$$\begin{aligned}
i\, c_j c_k &= 2i \sum_{p,q} \left([\underline{u}_p]_j \, f_p^\dagger + [\underline{u}_p]_j^* \, f_p \right) \left([\underline{u}_q]_k \, f_q^\dagger + [\underline{u}_q]_k^* \, f_q \right) \\
&= 2i \sum_{p,q} \Big([\underline{u}_p]_j [\underline{u}_q]_k \, f_p^\dagger f_q^\dagger + [\underline{u}_p]_j [\underline{u}_q]_k^* \, f_p^\dagger f_q + [\underline{u}_p]_j^* [\underline{u}_q]_k \, f_p f_q^\dagger \\
&\qquad\qquad + [\underline{u}_p]_j^* [\underline{u}_q]_k^* \, f_p f_q \Big) \\
&= 2i \sum_{p,q} \Big([\underline{u}_p \, (\underline{u}_q^*)^\dagger]_{jk} \, f_p^\dagger f_q^\dagger + [\underline{u}_p (\underline{u}_q)^\dagger]_{jk} \, f_p^\dagger f_q + [\underline{u}_p (\underline{u}_q)^\dagger]_{jk}^* \, f_p f_q^\dagger \\
&\qquad\qquad + [\underline{u}_p (\underline{u}_q^*)^\dagger]_{jk}^* \, f_p f_q \Big) \,.
\end{aligned} \tag{4.13}$$

4.2.5. Topological phases, the toric code and Ising anyons

There are two complementary perspectives on the Honeycomb model. For fixed vortex sector it is a model of free Majorana fermions, eq. (4.8). In some of these sectors the spectrum is gapless, but in others the fermions can be gapped and carry topological phases [51, 53, 54]. Introducing vortex defects may radically change the fermionic properties, for example by localising bound zero energy states at the vortex cores [40]. These bound fermionic states can, in turn, encode the exotic properties of anyons such as their braiding when the fermionic state is interpolated adiabatically between different vortex sector [52]. Alternatively, at low energies, the full spin model eq. (4.2) is a topologically ordered lattice model with two types of excitations: fermions and anyons. Depending on the parameter choices these anyons may be Abelian toric code anyons or non-Abelian Ising anyons.

Before moving on to outline the phase diagram of the model, we briefly introduce the notation we will use discuss the particles of the different anyon models [35]. (We will describe the fusion rules of the models, but not the braiding rules as we will not need them.) *Fusion rules* are written $a \times b = c$. This tells us that if particles a and b are brought close together then they act as if they were a particle of type c, when viewed on scales much larger than the separation between them. The particle types of the toric code are the vacuum 1, the anyons e and m and the fermion $\epsilon = e \times m$, which corresponds to fusing an e and m anyon together. The e, m and ϵ particles are all their own anti-particle, e.g. $e \times e = 1$. In the Ising anyon model, the particle types are the vacuum 1, the fermion ψ and the Ising anyon σ. The Ising anyons satisfy the non-Abelian fusion rule $\sigma \times \sigma = 1 + \psi$. Namely, if two σ are brought together they can either fuse to give the vacuum, 1, or a fermion, ψ. Again, the fermion is its own anti-particle $\psi \times \psi = 1$.

The phase diagram of the no-vortex sector [51] is shown in fig. 4.4. The dependence on J_x, J_y and J_z (all chosen positive and real) is drawn with K taken fixed and small. (It is also possible to drive phase transitions by tuning K to large values with the J's fixed.) The central point corresponds to $J_x = J_y = J_z$, and the corners of the diagram are the limiting cases of sending two of the J's to zero. There are four distinct phases, the B phase at the centre and three A phases that occurs in the corner regions. Taking the example

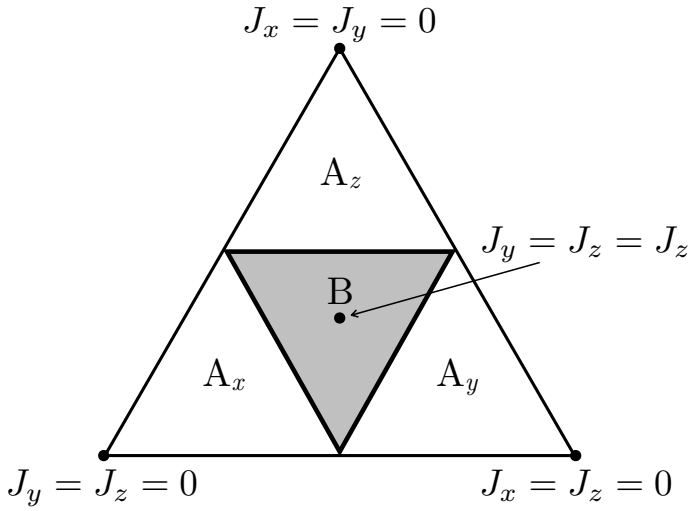


Figure 4.4. – Phase diagram of the no-vortex sector of the Kitaev honeycomb over parameters J_x, J_y and J_z. There are four distinct phases A_x, A_y, A_z and the central B phase. The boundary of the B phase with A_z occurs along the line $J_z = J_x + J_y$, and similar for A_x and A_y.

of the lower left corner, the boundary between A_x and B occurs when $J_x = J_y + J_z$.

The B phase carries a topological phase. However, while $K = 0$ the fermions are gapless, destroying the topological properties. Breaking time-reversal symmetry by setting $K > 0$ opens a gap in the spectrum, allowing the topological features to manifest. This is quantified by the bulk Chern number [41]. The Chern number can be written as an integral over individual spectal bands in momentum space. However, we will instead write it in terms of spatial two-point correlations (see Appendix C of [51]) as

$$\nu = \frac{3\pi i}{2} \sum_{i \in A} \sum_{j \in B} \sum_{k \in C} \left(\langle c_i c_j \rangle \langle c_j c_k \rangle \langle c_k c_i \rangle - \langle c_i c_k \rangle \langle c_k c_j \rangle \langle c_j c_i \rangle \right). \tag{4.14}$$

where A, B and C are subregions of the lattice such that A shares a boundary with both B and C, but all three only meet at a single point. Additionally, the three together should be enclosed within a topologically non-trivial loop. Finally, the ordering of A, B and C should be counter clockwise. The regions are sketched schematically in fig. 4.5. Using eq. (4.13) the Chern number can be evaluated for arbitrary states. (It is proven in [51] that in the translationally invariant case this expression is equivalent to the more conventional integral expressions in momentum space. In other cases it can be considered a generalisation of the translationally invariant case.) Evaluating ν for the fermionic ground

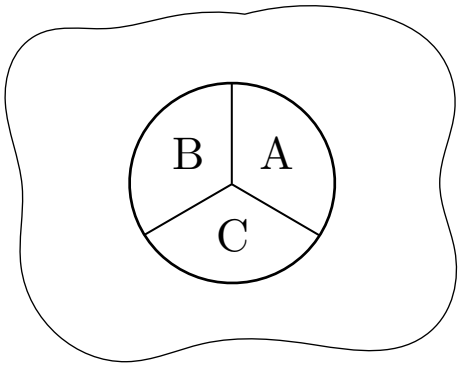


Figure 4.5. – Arrangement of regions A, B and C within the system (outer boundary) used to compute the Chern number from spatial two-point correlations, eq. (4.14). A, B and C only all meet at a single point, their arrangement is counter clockwise and combined together they are bounded by a topologically trivial curve.

state of the B phase using eq. (4.14) gives $\nu = \pm 1$. The Chern number encodes important information about the thermal transport (which we will return to in chapter 5) and the properties of vortex excitations. Vortex defects in the B phase bind Majorana zero-modes and behave (at large separations) like non-Abelian Ising anyons σ [52].

All of the A phases are gapped for all values of K. They do not carry topological phases. Their Chern number is trivial $\nu = 0$. Vortex defects in the A phases behave as Abelian anyons. In fact, Kitaev showed using perturbation theory (e.g. in the case $J_x \gg J_y, J_z$) that the A phases are each equivalent to the toric code [51]. The vortices of the A phases map to the e and m anyons of the toric code, though the details of this mapping is different for A_x, A_y and A_z. As an example, if we choose to make $J_z \gg J_x, J_y$ the vortices living on alternating rows of plaquettes are identified with e or m toric code anyons respectively [51]. (An eccentricity of the Kitaev honeycomb's representation of the toric code is that it supports states with a single e particle and a single m particle, which are not allowed in the stabiliser description of the toric code (chapter 2). We will explain how this is possible at the beginning of chapter 6.)

The picture described by the phase diagram breaks down for vortex sectors far from the ground state. Non-topological interactions between the vortices can drive the system into new regimes. Regular patterns of Ising vortices can produce novel emergent topological phases with different Chern numbers [53, 54]. A particular example of this is the full-vortex sector (all $W_p = -1$), which has $\nu = \pm 2$ [53]. Disordered configurations of Ising vortices at high enough density create a *thermal metal* [87]. This is a gapless phase with

a divergent density of states at zero energy [88, 89]. We will return to a discussion of this exotic metallic phase in chapter 6.

4.3. Summary

This chapter gives a minimal introduction to the deep and complex physics of the Kitaev Honeycomb, starting with its definition as a spin model and leading up to Kitaev's solution in terms of vortices and fermions. It ends with a brief discussion of the physical properties of the solution in different regimes. Among these properties is the presence of topological phases in the physics of the fermions. The next chapter returns to this topic to study edge currents in the topological phases of the honeycomb model.

To again highlight a particularly important aspect of the transformation from spins to fermions and vortices, the fermionic states must be projected into a particular subspace of their Hilbert space to match up to spin states. This issue will be returned to in chapter 6 in order to study the thermodynamics of the full spin model.

5.

Topological edge currents

In this chapter we study topological edge currents in Kitaev's honeycomb. This is a signature of the bulk-boundary correspondence exhibited by topological phases [51, 90]. Edge currents are a property of single vortex sectors, arising from the topological phases of the fermions in these sectors. Therefore, in this chapter we restrict our attention to the fermions living in certain special fixed vortex sectors. Additionally, since we are only interested in the physics of the topological phases we do not project the fermionic states into the spin subspace. The topological phases of Kitaev's honeycomb are topological superconductors [43].

Topological superconductors are extremely secretive. They do not conduct charge currents, they expel magnetic fields and their topological signatures are hidden from any local observable [43]. Through an effective gravitational description it has been shown that this "dark matter" of solid-state physics gives rise to a conformal field theory (CFT) at its boundary [90]. This CFT description of the topological edge states is remarkably robust. Unlike genuinely $1+1$-dimensional CFT, associated with fine-tuned critical points, these edge CFT typically persist across finite regions of the superconductor parameter space. This makes such systems an exciting medium for investigating direct signatures of conformal invariance both theoretically and experimentally.

There is a tight relationship between the Chern number ν, which describes the bulk physics, and the central charge c of the edge CFT. They are related by $c = \nu/2$ [90]. Conformal field theory predicts the topological states living on the edge will carry an energy current that scales as $I_{\mathrm{CFT}} = \pi/12\, k_B^2/h\, cT^2$ with temperature T, which must be small [91–93]. Here k_B is Boltzmann's constant and h is Planck's constant, these fix the units of I_{CFT} to be energy / time. In the rest of this chapter we will set k_B and h to one. Practically, however, the edge states are not perfectly isolated from the rest of the system. They have a finite penetration into the bulk, which has its own thermal behaviour. Inevitably these behaviours will mix and a natural question is whether it is still possible

to obtain conclusive signatures of CFT thermal properties.

Here we investigate the edge physics of a topological superconductor using Kitaev's honeycomb model as a microscopic description. We begin by considering the competing factors that would cause the behaviour to depart from the ideal CFT scaling as T becomes large and as $T \to 0$. This allows us to identify the range of temperatures for which the currents can be expected to obey the CFT prediction, section 5.1. Following this, in section 5.2 we construct an energy current operator and study its behaviour including its dependence on temperature and its topological robustness. We probe the edge currents in three different topological phases of the Kitaev honeycomb: the toric code phase with $\nu = 0$, the B of the no-vortex sector with $\nu = 1$ [51] and the topological phase of the full-vortex sector with $\nu = 2$ [53]. This section ends with a closer examination of one of the ways the currents depart from their idealised behaviour, which is the finite size scaling we observe at zero temperature.

5.1. Ideal temperature range

Consider the honeycomb lattice in a phase with non-trivial Chern number, wrapped around a cylinder with a perimeter of length L and height D. In this case the Hamiltonian can be solved using plane wave solutions in the periodic direction, which we chose to be the x direction. These solutions are given by $u_{k,d}(x, y) = v_d(y)\, e^{ipx}$ where $p = (2\pi/L)\, n$ with $n = 1, 2, \ldots, L$ is a momentum index in the x direction. The index $d = 1, 2, \ldots, D$ labels the different solutions $v_d(y)$ of the one dimensional problem in the open direction. The energies of these states are plotted in fig. 5.1. Both positive and negative energies are shown, however, as discussed in section 4.2.4 these do not all correspond to independent states. Positive energy states describe the possible quasiparticle excitations. The ground state corresponds to filling all of the negative energy states.

A non-zero Chern number indicates that the spectrum contains chiral midgap modes [41]. These modes have linear dispersion and are exponentially localized on the boundaries of the system, as shown in fig. 5.1. Thermal excitation of these modes give rise to an energy current on the edge. It is possible to heuristically estimate this current. At low

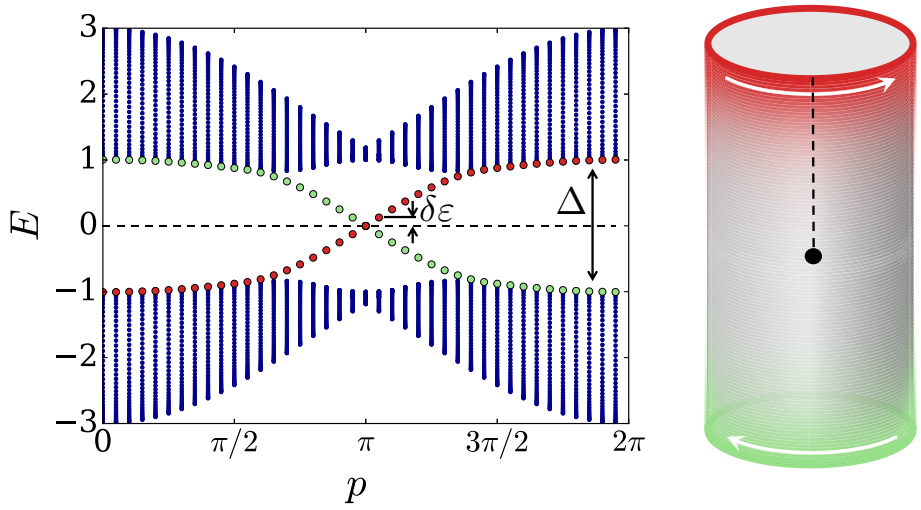


Figure 5.1. – Energy eigenvalues of the Hamiltonian (Left) for a cylinder of perimeter L and height D, as a function of momentum in the periodic direction $p = (2\pi/L)\, n$ for $n = 1, 2, \ldots, L$. Each value of p corresponds to D different energies, each associated to different solutions of the one dimensional problem in the open direction. The bulk energy gap, Δ, of the model and the discretisation gap, $\delta\varepsilon$, due to the length of the edge L are shown. Of particular interest are the edge states, which are highlighted in red and green. In the continuum limit $L \to \infty$ these would make up two midgap bands that cross at $p = \pi$. Close to zero energy states of the two bands are exponentially localised at opposite ends of the cylinder (Right). In contrast all bulk states, above the bulk gap Δ, vanish on the open boundaries. Chiral energy currents flow in opposite directions on the two edges are indicated by the arrows (Right).

temperatures T compared to the bulk energy gap Δ, the current is given by

$$I_{\text{edge}} = \int_{\varepsilon(p)\geq 0} n_\beta\big(\varepsilon(p)\big)\, \varepsilon(p)\, v(p)\, \frac{\mathrm{d}p}{2\pi} \approx \sum_{p:\, \varepsilon(p)\geq\delta\varepsilon}^{\varepsilon(p)\leq\Delta} n_\beta(\varepsilon)\, \varepsilon(p)\, \frac{\delta\varepsilon}{2\pi}\,, \tag{5.1}$$

where $n_\beta(\varepsilon) = 1/(1 + e^{\beta\varepsilon})$ gives the fermionic occupancies, $v(p) = \mathrm{d}\varepsilon/\mathrm{d}p$ is the velocity of the edge mode and $\varepsilon(p)$ is the edge mode's dispersion relation, shown in fig. 5.1 (Left). Intuitively, the expression $\varepsilon(p)\, v(p)$ combines the energy the edge mode is carrying with the rate at which it is being transported. Multiplying this with the probability the mode is occupied, $n_\beta(\varepsilon)$, then gives the expectation value of the amount of energy transported along the edge per unit time. The linear dispersion of the edge mode means energy is transported in one direction along the edge (i.e. the current is chiral) and, since the dispersion of the two edges is in opposite direction, energy flows in opposite directions on the two boundaries. This is indicated by the arrows on the right panel of fig. 5.1.

In the limit of infinite size the edge modes of the cylinder are gapless around $p = \pi$. A finite circumference L induces an infra-red cut-off given by $\delta\varepsilon \sim 1/L$ due to the discreteness

of the momentum values $p_{\cdot\cdot}$. Additionally, the bulk gap Δ naturally sets an ultra-violet cut-off. (Δ is evaluated by imposing periodic boundary conditions in both directions.) These conditions determine the limits of the sum eq. (5.1). If we send $\delta\varepsilon \to 0$ and $\Delta \to \infty$ then the integral can easily be computed and eq. (5.1) evaluates to the CFT current, $I_{\mathrm{CFT}} = {}^{\pi}/_{12}\, c\, T^2$, exactly. However, these limits place bounds on the range of temperatures at which we expect to see CFT current behaviour, given by

$$L^{-1} \ll T \ll \Delta. \tag{5.2}$$

Furthermore, eq. (5.1) assumes perfect distinguishability of the edge modes. While this is energetically possible, it cannot be achieved by local position measurements of the current. Indeed, while edge modes are exponentially localized they still have a finite penetration into the bulk of the system, as shown in fig. 5.1 (Right). Hence any attempt to probe them theoretically or experimentally from a microscopic model needs to consider current contributions from all states, not just the midgap modes [51].

5.2. Energy Currents

In section 5.2.1 we construct an energy current operator and express it in terms of two-point fermionic correlators. Section 5.2.2 shows the behaviour of the thermal currents $I(T)$ for various topological phases. The currents are seen to yield the central charges. Additionally, we show that when the system parameters are varied the currents can be used to cleanly signal topological phase transitions at finite temperature. In section 5.2.3 we show the topological nature of the central charge by studying its resilience to random disorder and boundary defects. Finally in section 5.2.4 we address the currents at $T = 0$.

5.2.1. Defining a current operator

As outlined in section 4.2, for fixed vortex sector the spin system reduces to one of free Majorana fermions whose Hamiltonian $\overline{H}$ is given by

$$\overline{H} = \sum_j h_j \qquad \text{where} \qquad h_j = \frac{i}{4} \sum_{k=1}^{N} A_{jk} c_j c_k \,, \tag{5.3}$$

where $A_{jk} = \pm J_\alpha$ when j and k are nearest neighbours, $A_{jk} = \pm K$ when they are next-nearest neighbours and $A_{jk} = 0$ otherwise. We have separated the Hamiltonian into a sum of terms h_j centred on site j. Although the sum defining h_j is over the whole lattice, the couplings A_{jk} are short range and so these terms have localised support about site j. An energy current operator is defined through the Heisenberg equation for the local energy terms, given by

$$\frac{\mathrm{d}}{\mathrm{d}t} h_j \;=\; -i[\,\overline{H}, h_j\,] \;=\; -i\sum_k [h_k, h_j] \;\equiv\; \sum_k I_{jk}\,, \tag{5.4}$$

where the final step defines the current operator $I_{jk} \equiv -i[h_k, h_j]$. This is the differential form of a lattice continuity equation for energy, linking the rate of change of energy at one lattice site to the current it exchanges with other lattice sites.

We next expand the current operator in terms of the lattice fermions c_i to obtain an expression for I_{jk} based on two-point correlators. It is useful to break up the local energy terms h_j and h_k in the following way

$$h_k \;=\; \frac{i}{4}\sum_l A_{kl} c_k c_l \quad = \quad \frac{i}{4}\left(A_{kj} c_k c_j + \sum_{l\neq j} A_{kl} c_k c_l \right)\,,$$

with an analogous expression for h_j. When expanding the commutator, the terms then group together so that it is easy to identify which vanish. The current operator becomes

$$\begin{aligned} I_{jk} \;=\; \frac{i}{16}\Bigg(& A_{kj}A_{jk}\,[c_k c_j, c_j c_k] + \sum_{m\neq k} A_{kj}A_{jm}\,[c_k c_j, c_j c_m] + \sum_{l\neq j} A_{kl}A_{jk}\,[c_k c_l, c_j c_k] \\ & + \sum_{l\neq j} A_{kl}A_{jl}\,[c_k c_l, c_j c_l] + \sum_{l\neq j}\sum_{m\neq k,l} A_{kl}A_{jm}\,[c_k c_l, c_j c_m] \Bigg)\,. \end{aligned}$$

The first term vanishes since $c_k c_j c_j c_k = c_j c_k c_k c_j = 1$. Similarly, we can show that since the four indices in the last term are different that commutator also vanishes. Manipulating the other terms yields

$$I_{jk} \;=\; \frac{i}{8}\left(\sum_{m\neq k} A_{kj}A_{jm}\, c_k c_m + \sum_{l\neq j} A_{kl}A_{jk}\, c_l c_j - \sum_{l\neq j,k} A_{kl}A_{jl}\, c_k c_j \right)\,. \tag{5.5}$$

Substituting the expression for $ic_j c_k$ from section 4.2.4 we can write the current operator

in terms of the energy eigenmodes. Expectation values of the current operator can easily be obtained in the energy basis. The non-zero elements are associated with the particle and hole occupations given by,

$$\langle f_p^\dagger f_q \rangle = {}^1\!/_{\overline{\mathcal{Z}}}\ \mathrm{tr}\Big(e^{-\beta \overline{H}} f_p^\dagger f_q \Big) = \frac{1}{e^{\beta\varepsilon_p}+1}\,\delta_{pq}$$
$$\langle f_p f_q^\dagger \rangle = {}^1\!/_{\overline{\mathcal{Z}}}\ \mathrm{tr}\Big(e^{-\beta \overline{H}} f_p f_q^\dagger \Big) = \frac{e^{\beta\varepsilon_p}}{e^{\beta\varepsilon_p}+1}\,\delta_{pq}\,,$$

where $\overline{\mathcal{Z}} = \mathrm{tr}(e^{-\beta\overline{H}})$ is the partition function. (A *vectorised* expression for the current operators which is faster to evaluate is given in note B.1.5.)

The edge current is computed as the total net energy flux on an edge, we first compute the net energy flux around the cylinder as a function of height y. Then we sum up all these local currents between the middle of the system and the boundary, as shown with the dashed line on fig. 5.1 (Right), to obtain

$$I(T) = \sum_{y=D/2}^{D} \left(\sum_{\langle j,k\rangle\,:\,y} \langle I_{jk}(T)\rangle \right) . \tag{5.6}$$

The inner sum is performed over links $\langle j,k\rangle$ that cross the cut at height y, while the outer sum captures the current on one edge only.

5.2.2. Central charges

We are now in a position to study the behaviour of the energy currents $I(T)$ as we vary the temperature T. Before we can compare the scaling of $I(T)$ with the CFT prediction, I_{CFT}, there are two aspects of the energy current we need to address. These are the contribution from the bulk and the bounds set by condition (5.2).

The prediction of CFT is that the currents should vanish at $T = 0$ in the thermodynamic limit, $L \to \infty$. By carrying out a finite size scaling analysis of our data $I(T = 0)$ we find that this is not the case. We identify a non-zero value $I_0 = I(T = 0, L \to \infty)$ at zero-temperature, as shown in fig. 5.2 (Left), the value of which depends on the microscopic parameters of the model. In order to compare the energy currents of different phases we subtract the I_0 contribution. This is equivalent to considering the conductivity (the rate

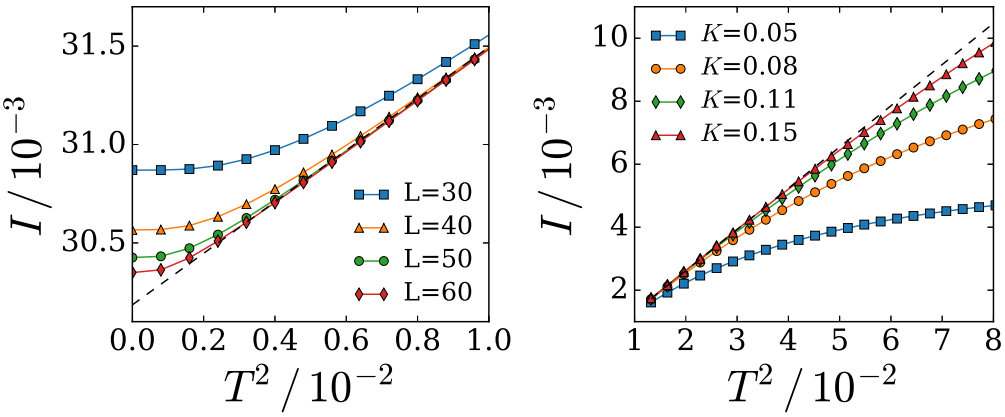


Figure 5.2. – Low and high T behaviour of $I(T)$. (Left) In the low-T case increasing system size, L, recovers the T^2 behaviour at lower temperatures. (Right) In the high-T case increasing $K \propto \Delta$ delays the divergence of the curves. The (Left) panel includes the I_0 offset, whereas we have removed it from the (Right) panel. In both plots $J = 1$. In the (Left) plot $K = 0.15$ and in the (Right) plot $L = 60$.

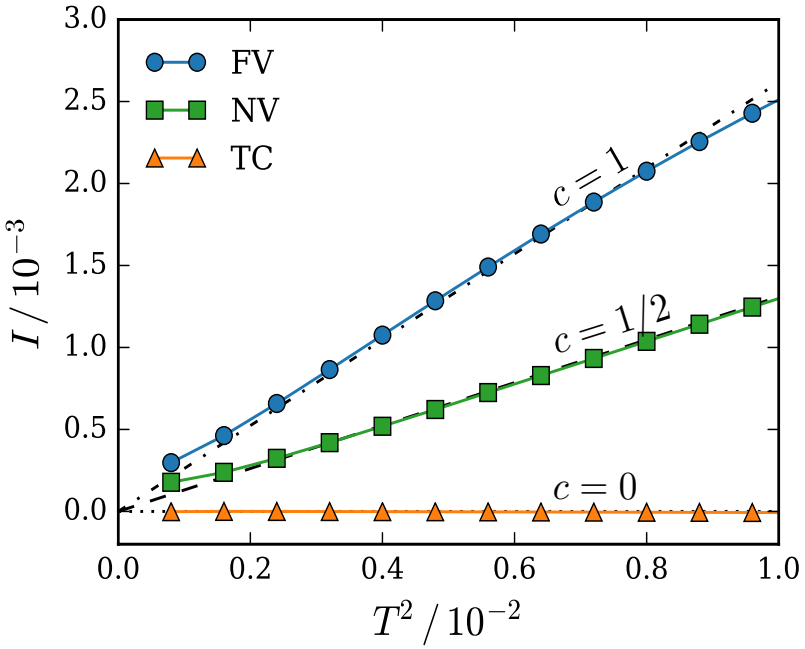


Figure 5.3. – Edge current for the no-vortex sector (NV), full-vortex sector (FV) and Toric Code phases (TC) as a function of T^2. The currents are shifted to pass through the origin, clearly revealing the T^2 scaling. The central charges correspond to dotted line for $c = 0$, dashed line for $c = 1/2$ and dashed-doted line for $c = 1$. In these plots we set $K = 0.15$. For the NV and FV plots we set all $J = 1$, and for the TC plot we take $J = 0.1$. Each is plotted for system size $L = 60$.

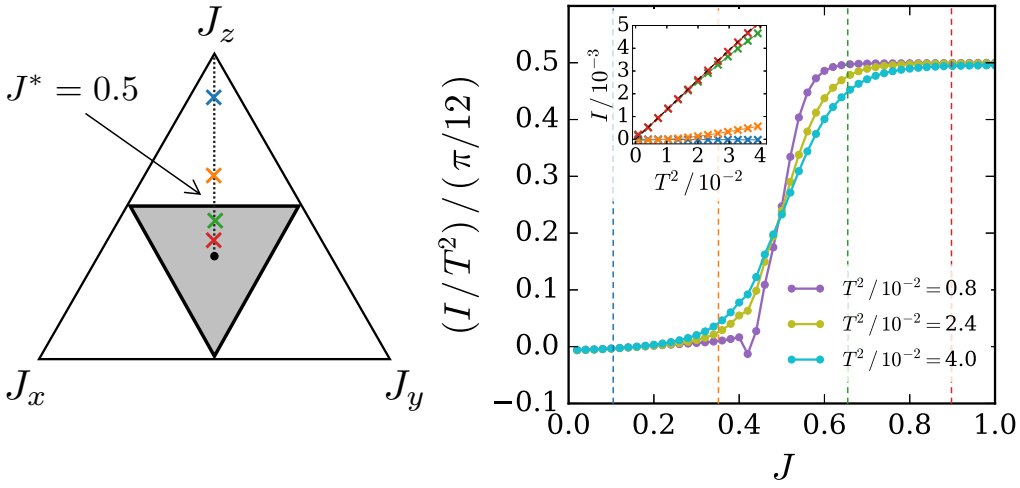


Figure 5.4. – Edge currents indicate a transition between different topological phases at finite temperature. Here we vary $J_x = J_y = J$ for fixed $J_z = 1$, corresponding to the path through parameter space shown on the triangle. Scaled edge currents I/T^2 are plotted against J for different values of T. We see a jump through the phase transition that sharpens with decreasing T. (Inset) I against T^2 plotted for the values of J indicated on the triangle. The currents are seen to be robust in the topological phase (red and green crosses) and vanish in the Toric code phase (orange and blue crosses). These plots are for $L = 52$ and $K = 0.15$.

of change of current with T) rather than the net current value.

Condition (5.2) puts bounds on the temperature range we expect to see CFT scaling. In fig. 5.2 we plot the currents at the low-T and high-T ends of these bounds, for the no-vortex sector. We see the finite size dependence of the low-T bound. In the low-T plot we have not removed the zero temperature value, meaning as L is increased the currents at $T = 0$ tend to I_0. This generic low-T behaviour is seen for all model parameters. In contrast at large T the currents are sensitive to the bulk gap Δ but independent of system size. We observe that the temperature at which I starts to clearly deviate from T^2 scaling (greater than a 10% difference) grows linearly with Δ. To summarise, we have verified that the temperature range of interest has lower and upper bounds that scale proportional to the inverse system size and the energy gap respectively.

For temperatures inside the bounds of condition (5.2) we find energy currents that scale as CFT predicts. To demonstrate this we plot $I(T)$ against T^2 for the Toric code, no-vortex and full-vortex phases in fig. 5.3. We find the currents are in excellent agreement with the CFT predictions where for TC, $c = 0$, for NV, $c = 1/2$ and for FV, $c = 1$.

Since the central charge only takes rational values [94] the current $I(T)$ should jump

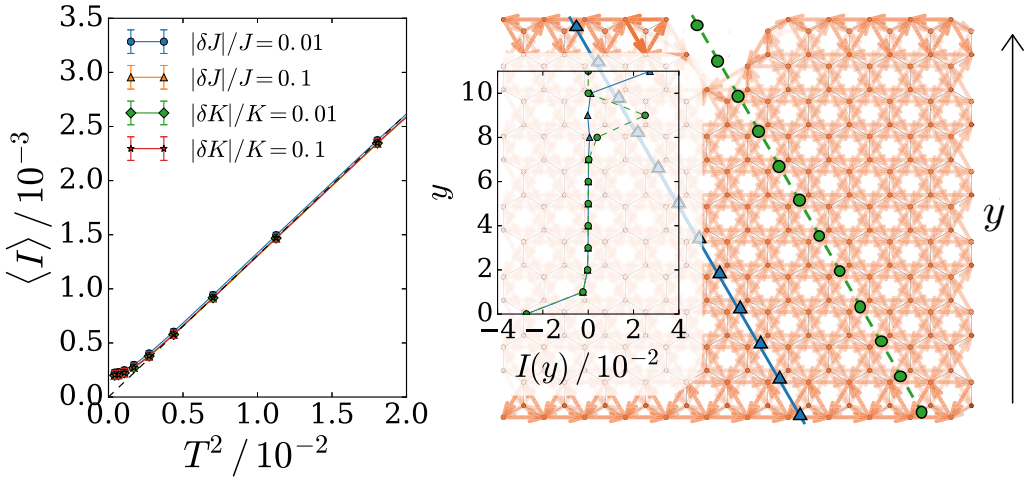


Figure 5.5. – Robustness of the edge currents to J and K disorder (Left) and lattice boundary defects (Right). In the (Left) plot we add a random component to either every J term in the Hamiltonian or every K term. We plot curves for $|\delta J|/J$ and $|\delta K|/K = 0.01$ and 0.1, setting $J = 1$ and $K = 0.1$, for $L = 40$. Each point is averaged over 100 disorder realizations. In both cases the energy currents remain unaffected.

as we move between two phases with different Chern numbers. This gives $I(T)$ the useful property that it identifies topological phase transitions at finite temperatures. In fig. 5.4 (Right) we plot the energy currents as we transition from the TC to NV phases at different temperatures. The specific transition we probe is achieved by tuning the J couplings, and is illustrated by the parameter-space diagram on the left of fig. 5.4. Higher temperatures are seen to smear out the transition and at the critical point we find a crossing between the different temperature curves. As the finite temperature behaviour of the energy currents is determined by the CFT at $T = 0$ the currents are a useful tool to characterize topological phases at finite temperature.

5.2.3. Topological robustness

The origin of the energy current $I(T)$ is topological, so we expect it to be robust against local perturbations to the Hamiltonian. We have investigated this property in two settings. First we introduced disorder to the couplings J and K. We begin with the no-vortex sector and add a random component to either every J or K value so that *e.g.* $J \to J + \delta J$ where δJ is a uniform random number between $-|\delta J|$ and $|\delta J|$. The currents are averaged over disorder $\langle I(T) \rangle$ and plotted as a function of T^2 for disorder strengths $|\delta J|/J$ and $|\delta K|/K$ of 1% and 10%, as shown in fig. 5.5 (Left). We see no impact of this disorder on the

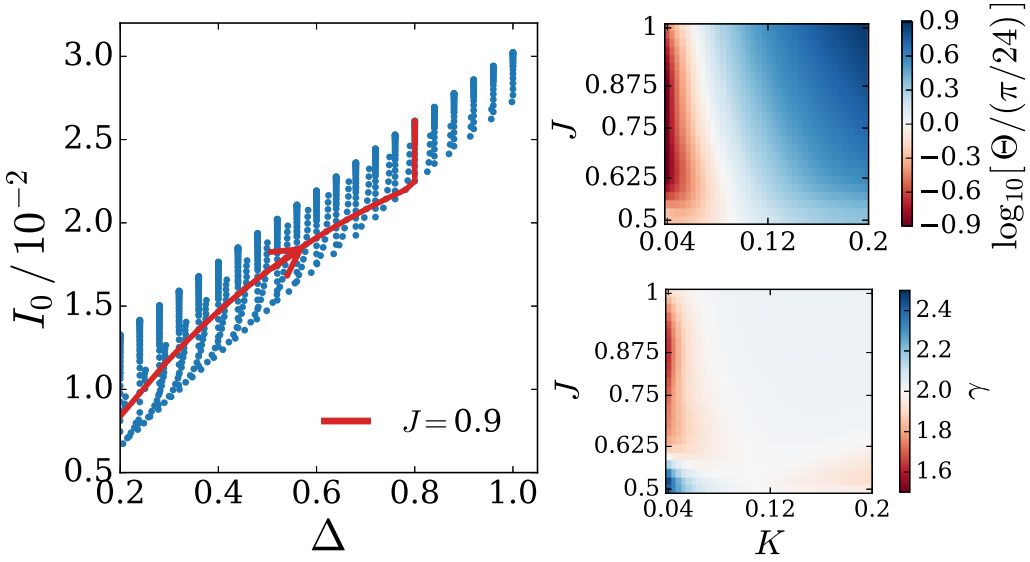


Figure 5.6. – Finite size scaling analysis of the zero-temperature current in the no-vortex sector. This is carried out by fitting values $I(0, L)$ for $L = 10, 12, \ldots, 52$ to eq. (5.7), while varying J and K within the no-vortex phase with $J_z = 1$. (Left) We plot the extrapolated values I_0 as a function of the fermion gap Δ. The relationship is non-unique. Onto the plot we have added the path of a single J while K increases. (Right) The pre-factor Θ and scaling exponent γ are plotted over J and K. We see for most parameters $\gamma \approx 2$. Interestingly, Θ is not found to have its CFT predicted value $\pi/12\,c$ (where $c = 1/2$) over most of the range investigated.

currents. Additionally, we introduce a boundary defect by removing sites from the edge. The energy currents near this defect are plotted in fig. 5.5 (Right). We find the impact of the defect is to divert the edge current around the missing lattice sites. A current of equal intensity flows along the new edge.

5.2.4. Currents at $T = 0$

We now give a fuller account of the currents at zero temperature, $I(T = 0)$. We previously observed that at zero-temperature the current includes strong finite size effects as well as an offset I_0. These different elements are investigated by fitting values $I(T = 0, L)$ computed for different system sizes L to a scaling form

$$I(T = 0, L) = I_0 + \frac{\Theta}{L^\gamma}\,. \tag{5.7}$$

Finite size scaling is also found in conformal field theory. At $T = 0$ CFT predicts that the currents should depend on the circumference of the edge L through the scaling relation $I(L) = (\pi/12)\,c\,L^{-2}$. (The energy current is predicted to be the $\tau - x$ component of the

energy-momentum tensor. An elementary conformal mapping shows that on a cylinder $\langle T_{\tau x} \rangle = \pi c/12L^2$ [94].)

We plot the values obtained for γ, Θ and I_0 over a range of system parameters in the no-vortex phase in fig. 5.6. The fits are obtained by computing $I(T, L)$ at $T = 10^{-6}$ for system sizes $L = 10, 12, \ldots, 52$ and fitting to eq. (5.7). We vary J and K and find good agreement of γ with the CFT exponent $\gamma \approx 2$ for most values of the couplings. The value of γ depart from 2 at small K and small J and we attribute this to quasi-critical effects as the gap becomes small [95]. The extrapolated values I_0 show a complicated dependence on the system parameters. In particular we note that by varying K for constant J we can alter the value of I_0 without changing the gap Δ, as shown by the red trajectory plotted on fig. 5.6 (Left). In addition, the pre-factor Θ does not correspond to the CFT prediction $\Theta = (\pi/12)\, c$ over most of the range of J and K. Hence, I_0 and Θ are strongly dependent on the microscopic parameters of the model, unlike the predictions of CFT. This demonstrates that, due to their chiral nature, the bulk states also contribute significantly to the energy current.

5.3. Conclusions

In conclusion we have presented a method to obtain the energy currents of a topological phase from a microscopic model. We have studied the topological properties of these currents at finite temperature and the corresponding transitions between different topological phases. Our results confirm predictions made using an effective description in terms of conformal field theory [51].

Such CFT descriptions have been studied extensively at both the effective and microscopic level in the context of edge states of the quantum Hall effect [96, 97], but there is a major distinction between that setting and ours. Here, we have focused on the thermal transport appropriate to a topological superconductor as opposed to the charge currents found in the quantum Hall effect.

We have seen that due to their CFT origin the energy currents are able to discriminate clearly between different topological phases at finite temperature. Moreover, we have demonstrated that these currents are very robust. Unlike the fragile physics of criticality

that typically has zero measure in the parameter space, the central charges evaluated through edge energy currents are robust against significant variations of the Hamiltonian parameters. We demonstrated that, due to their topological origin, these currents are also robust against bulk disorder or even the introduction of boundary defects. This resilience makes them the ideal method to experimentally probe topological superconducting phases and reveal their conformal behaviour.

6.

Kitaev's honeycomb at finite temperature

This chapter studies the Kitaev honeycomb spin model at finite temperature by exploiting its exact solution in terms of fermions and vortices. Kitaev's solution of the honeycomb model separates the spin degrees of freedom into fermions and a gauge field whose physical excitations, the Wilson loops, are the vortices. However, the two are not independent. The state of the gauge field determines the fermionic spectrum, but the connection goes even deeper. The states of the fermions must be projected into a particular subspace in order correspond to spin states and, as shown by [85], this projection further couples the state of the gauge field with the allowed total parity of the fermions. The consequences of this state level coupling between the gauge field and the fermions have never been fully explored. In this chapter we will first investigate some of the effects this has. In particular we will argue that it plays an important role in enforcing the fusion rules of the quasiparticle excitations and that it determines the topological degeneracy of the ground states of the model. The projection of the fermions is extended to finite temperatures and combined with a Monte Carlo procedure that samples from finite temperature gauge configurations in order to study the finite temperature properties of the spin model.

The finite temperature physics of the model has previously been studied by other authors [98, 99]. However, these studies have neglected some of the finer points of the model. In particular the projection of the fermionic states, which we show at the beginning of this chapter gives rise to many of the topological properties of the model. By incorporating these elements we are able to identify some of the topological properties of the model, such as its topological ground state degeneracy, in its thermodynamics. Following this, we study one of the interesting high temperature behaviours of the model that has not received the close attention it deserves before. Namely, the emergence of a thermal metal (an exotic heat conducting phase) from the Ising anyon phase at high temperature.

The chapter is organised as follows. A procedure to compute projected finite temperature expectation values for the fermions is described in section 6.1. During this section

we discuss some of the consequences of applying the projector. These consequences are further explored in section 6.2 where we consider the ground states of the model and show a link between the topological degeneracy of the ground states and the fermionic projection. Section 6.3 outlines the Monte Carlo approach we use to study the finite temperature behaviour of the full model. Finally, our results for the spin model thermodynamics are discussed in section 6.4.

6.1. Thermal states of the fermions

In this section we study the thermal states of the fermionic degrees of freedom. After a transformation from the original spin-1/2 particles, the fermions are described by superconducting Hamiltonian $\overline{H}$. However, the cost of obtaining this simplified representation is the need to apply a projector $\mathcal{P}$ to the fermionic states in order to obtain *physical* (spin) states. The physical projector is given by

$$\mathcal{P} = \prod_i \frac{\mathbb{1} + D_i}{2} = \mathcal{S}\,\mathcal{P}_0\,, \tag{6.1}$$

where the second equality was proven by Pedrocchi et al. [85]. Crucially, they showed that $\mathcal{P}$ divides up such that $\mathcal{S}$ acts purely on the gauge degrees of freedom u_{ij} and $\mathcal{P}_0$ fixes a physical parity for the allowed occupancies of the fermion eigenmodes. The projector $\mathcal{S}$ puts the gauge field into a superposition of all possible (topologically trivial) string configurations consistent with the vortex pattern, and for a square system $\mathcal{P}_0$ is given by

$$\mathcal{P}_0 = \frac{1}{2}\left(1 + \det(Q)\Big(\prod_{(i,j)} u_{ij}\Big)\mathcal{F}\right)\,. \tag{6.2}$$

In this case, the i of the (i,j) pair in the product over u_{ij} is chosen to always be the member of the black sublattice as drawn in fig. 4.3. The orthogonal matrix Q is the matrix that brings the $\overline{H}$ into Majorana normal form, eq. (4.10). Operator $\mathcal{F}$ counts the parity of the occupancies of the diagonal fermions f_p, $f_p^\dagger$ and is given by

$$\mathcal{F} = \prod_p (1 - 2f_p^\dagger f_p)$$

The eigenvalues of $\mathcal{F}$ are +1 (when an even number of modes are excited) and -1 (when the number is odd). The projector $\mathcal{P}_0$ therefore projects either onto even or odd fermion parity depending on the gauge configuration, encoded in the value of $\det(Q)\prod_{(i,j)} u_{ij}$. As noted in [85], importantly this factor is invariant under local gauge transformations D_i, which is the same as saying that $\mathcal{P}_0$ commutes with $\mathcal{S}$.

This coupling of the gauge field and the total fermion parity plays an interesting role in the emergent physics of the model. In particular, it enforces the anyonic fusion rules and sets the topological degeneracies (which we will consider in detail in section 6.2). The clearest example of this is found in the fusion rules of the toric code. Recall that the toric code has three particle types e, m and ϵ, which are all their own antiparticles. All allowed particle configurations should be connected to vacuum, meaning if we fuse all the particles present nothing remains. In the honeycomb model this condition is enforced by the fermion parity projection. The honeycomb vortices correspond to e and m anyons and the fermionic quasiparticles to ϵ. An even fermion parity implies that only states containing even numbers of ϵ quasiparticles are permitted, whereas odd parity means there is a residual ϵ particle left over when all the other fermions are annihilated together. For both the no vortex sector and sectors that contain pairs of e (or m) type vortices we find that the fermion parity projects to even parity, which correctly restricts us to states with total vacuum charge. Whereas, considering sectors that contain odd numbers of e and m vortices we find the fermion parity projects to odd parity. In this slightly odd way the net particle charge in the system remains vacuum, $\epsilon \times e \times m = \epsilon \times \epsilon = 1$, even though the two fermions (the quasiparticle fermion and the $e \times m$ fermion) manifest in entirely different degrees of freedom.

Consider computing the expectation value of a spin observable $\mathcal{O}$ for the projected physical state $|\Psi_{\text{spin}}\rangle = \mathcal{S}\,\mathcal{P}_0\,|\Psi_{u_{ij}}\rangle$, where $|\Psi_{u_{ij}}\rangle$ is a fermionic state with support in a single vortex sector, calculated from some fixed pattern of u_{ij}. The projector $\mathcal{S}$ decomposes into a sum over products of D_i. We can therefore address $\mathcal{S}$ by showing that

$$[\overline{\mathcal{O}}, D_i] = 0 \quad \text{for all } i\,, \tag{6.3}$$

where $\overline{\mathcal{O}}$ is the spin observable $\mathcal{O}$ translated into the fermion picture. This should be the

case for all physical quantities as the u_{ij} are gauge degrees of freedom.

Whereas $\mathcal{S}$ can be handled just by studying the operator $\mathcal{O}$, the part $\mathcal{P}_0$ unavoidably depends on the state $\left|\Psi_{u_{ij}}\right\rangle$. More generally, we will want to consider mixed states described in the fermionic picture by density operator $\hat{\rho}$. Assuming that the observable obeys condition (6.3), the projected expectation value is given by

$$\langle\mathcal{O}\rangle = \mathrm{tr}(\overline{O}\,\mathcal{S}\mathcal{P}_0\,\hat{\rho}\,\mathcal{S}\mathcal{P}_0) = \mathrm{tr}(\overline{\mathcal{O}}\,\hat{\rho}\,\mathcal{S}\mathcal{P}_0)\,.$$

We have used the that fact that if an operator $\overline{\mathcal{O}}$ commutes with $\mathcal{S}$ then it also commutes with $\mathcal{P}_0$. This can be seen to be true by considering the original expression for the full physical projector $\mathcal{P} = \prod_i {}^1\!/_2(\mathbb{1} + D_i)$, from this it is clear that if an observable meets condition (6.3) then it commutes with the full $\mathcal{P}$. For simple states, such as the thermal state, the density operator is diagonal in the energy eigenbasis. Additionally, projector $\mathcal{P}_0$ is diagonal in this basis. Therefore, we can complete the trace in the basis of quasiparticle excitations to give

$$\langle\mathcal{O}\rangle = \sum_{\{\underline{n}\}} p(\underline{n})\,\overline{\mathcal{O}}(\underline{n})\,\Pi(\underline{n})\,, \tag{6.4}$$

where the vector $\underline{n} = (n_1, n_2, \ldots)$ gives the quasiparticle occupancies $n_q = 0, 1$ of the modes q. The probability weight of pure state $\underline{n}$ in the mixed state is given by $p(\underline{n})$. The quantity $\overline{\mathcal{O}}(\underline{n})$ is the observable evaluated for $\underline{n}$. Lastly $\Pi(\underline{n}) = 0$ or 1 kills terms that do not obey the parity condition.

Similarly we can use this approach to study the projected partition function of the fermions. The partition function encodes the thermodynamics of the fermions and normalises the thermal state probabilities $p(\underline{n})$. We can write it as

$$\overline{\mathcal{Z}}(\beta) = \mathrm{tr}\left(\mathcal{P}e^{-\beta\overline{H}}\right) = \sum_{\{\underline{n}\}} \Pi(\underline{n})\,e^{-\beta\mathcal{E}(\underline{n})} \tag{6.5}$$

where the energy associated with state $\underline{n}$ is given by $\mathcal{E}(\underline{n}) = \sum_p \frac{\varepsilon_p}{2}\,(2n_p - 1)$. Without projection, when $\Pi(\underline{n}) = 1$ for all $\underline{n}$, the partition function can be evaluated analytically and exact expressions for the thermodynamic properties of the system can be derived. These standard results are reviewed in section 6.1.1. When including the projection it is not possible to obtain closed expressions for the thermodynamic quantities. However, it

is still possible to compute them exactly. Section 6.1.2 describes an iterative procedure to do this. Further, we can show that this full treatment is not always necessary as in some cases many thermal quantities are identical in the projected and unprojected cases. This is shown to be true for systems with a zero energy fermion mode in section 6.1.3.

6.1.1. Unprojected fermions

Without projection, the occupancy of each mode p is statistically independent and the partition sum eq. (6.5) factorises over the modes to give

$$\overline{\mathcal{Z}} \;=\; \prod_p \left(e^{\beta\varepsilon_p/2} + e^{-\beta\varepsilon_p/2} \right) \;=\; \prod_p 2\cosh(\,{}^{\beta\varepsilon_p}\!/_2)\,.$$

This exact form can be used together with standard relations from statistical physics to derive expression for the thermodynamic quantities of interest, which are then evaluated for the numerically obtained spectrum of $\overline{H}$. Some of the most important of these are:

Free Energy

The free energy, $\overline{F} = \overline{U} - T\overline{S}$ where $\overline{U}$ is the internal energy and $\overline{S}$ the thermal entropy, connects the statistical description $\overline{\mathcal{Z}}$ to other thermodynamic quantities

$$\overline{F} = -T \ln\left(\,\overline{\mathcal{Z}}\,\right) = -T \sum_p \ln\Big[\, 2\cosh({}^{\beta\varepsilon_p}\!/_2) \,\Big]\;.$$

Internal Energy

The internal energy at thermal equilibrium $\overline{U} = \langle \overline{H} \rangle$ can be calculated from $\overline{F}$,

$$\overline{U} = \frac{\partial}{\partial\beta}(\beta\overline{F}) = -\sum_p \frac{\varepsilon_p}{2} \tanh\left(\frac{\beta\varepsilon_p}{2}\right)\;.$$

Heat capacity

The heat capacity describes the heat that the system needs to absorb to cause unit change in temperature,

$$\overline{C} = T\,\frac{\partial \overline{S}}{\partial T} = -T\,\frac{\partial^2 \overline{F}}{\partial T^2} = \sum_p \left(\frac{\beta\varepsilon_p}{2}\right)^2 \frac{1}{\cosh^2\left(\beta\varepsilon_p/2\right)}\,.$$

Quasiparticle occupations

The particle and hole occupations of the quasiparticle states obey fermionic statistics,

$$\langle f_p^\dagger f_q\rangle = \frac{1}{\mathcal{Z}}\,\mathrm{tr}\Big(\,e^{-\beta\overline{H}} f_p^\dagger f_q\,\Big) = \frac{1}{e^{\beta\varepsilon_p}+1}\,\delta_{pq} \tag{6.6}$$

$$\langle f_p f_q^\dagger\rangle = \frac{1}{\mathcal{Z}}\,\mathrm{tr}\Big(\,e^{-\beta\overline{H}} f_p f_q^\dagger\,\Big) = \frac{e^{\beta\varepsilon_p}}{e^{\beta\varepsilon_p}+1}\,\delta_{pq}\,. \tag{6.7}$$

Correlation functions

Using equations (6.6) and (6.7) we can write thermal expectation values for (ic_jc_k) from the expression obtained in section 4.2.4. We additionally exploit a property of the eigenvectors of the A matrix to simplify the final expression further: $[\underline{u}_p(\underline{u}_p)^\dagger]^* = -[\underline{u}_p(\underline{u}_p)^\dagger]$. This must hold in order for (ic_jc_k) to be a hermitian operator. The correlations are given by

$$\langle ic_jc_k\rangle = -i\sum_p [\underline{u}_p(\underline{u}_p)^\dagger]_{jk}\,\tanh\left(\frac{\beta\varepsilon_p}{2}\right)\,.$$

The notation here corresponds to the notation introduced in section 4.2.4.

6.1.2. Projected fermions

Evaluating the projected partition function is more complicated since $\Pi(\underline{n})$ depends on the total parity of the occupancies $\underline{n}$ and hence the modes are no longer independent. This section describes an iterative procedure that efficiently computes eq. (6.5).

Starting with a single quasiparticle mode n_1 the even $n_1 = 0$ and odd $n_1 = 1$ states are clear. The single-mode partition functions restricted to the even subspace $\mathcal{Z}_1^{\mathrm{e}}$ and the odd

subspace $\mathcal{Z}_1^{\mathrm{o}}$ are then simply given by

$$\mathcal{Z}_1^{\mathrm{e}} = e^{+\beta\varepsilon_1/2} \quad \text{and} \quad \mathcal{Z}_1^{\mathrm{o}} = e^{-\beta\varepsilon_1/2}\,.$$

Adding a second mode the even and odd subspaces of the two are again clear. The even states are $(0,0)$ & $(1,1)$, and the odd states are $(1,0)$ & $(0,1)$. Carrying forward the notation used above, we can write the two-mode partition functions of the subspaces as

$$\begin{aligned}\mathcal{Z}_2^{\mathrm{e}} &= \mathcal{Z}_1^{\mathrm{e}}\, e^{+\beta\varepsilon_2/2} + \mathcal{Z}_1^{\mathrm{o}}\, e^{-\beta\varepsilon_2/2}\,,\\ \mathcal{Z}_2^{\mathrm{o}} &= \mathcal{Z}_1^{\mathrm{o}}\, e^{+\beta\varepsilon_2/2} + \mathcal{Z}_1^{\mathrm{e}}\, e^{-\beta\varepsilon_2/2}\,.\end{aligned}$$

As we continue to add further modes the number of states grows exponentially and it quickly becomes impractical to track all even and odd states. However, these expressions for the two-mode partition function have compressed the information about their occupations into just their total parity. To add the next mode we do not need to track the exact state vectors, instead we can proceed by adding the even and odd states of n_3 to the even and odd subspaces of the combined n_1 and n_2. This gives an iterative scheme for evaluating the even and odd restrictions of the partition function

$$\mathcal{Z}_n^{\mathrm{e}} = \mathcal{Z}_{n-1}^{\mathrm{e}}\, e^{+\beta\varepsilon_n/2} + \mathcal{Z}_{n-1}^{\mathrm{o}}\, e^{-\beta\varepsilon_n/2}\,, \tag{6.8}$$

$$\mathcal{Z}_n^{\mathrm{o}} = \mathcal{Z}_{n-1}^{\mathrm{o}}\, e^{+\beta\varepsilon_n/2} + \mathcal{Z}_{n-1}^{\mathrm{e}}\, e^{-\beta\varepsilon_n/2}\,. \tag{6.9}$$

(Evaluating equations (6.8) and (6.9) could potentially be numerically unstable, see note B.1.3 for a discussion of how we avoid this.)

Thermodynamic quantities can be computed from the partition function. However, this involves differentiating $\mathcal{Z}$, and so would require evaluating $\mathcal{Z}$ multiple times at different temperatures. Additionally, quantities such as the correlations $\langle ic_jc_k\rangle$ cannot be computed in this way. Instead, thermal expectation values can be computed from their statistical expressions. Typically this will involve the quasiparticle occupancies $\langle n_p\rangle = \langle f_p^\dagger f_p\rangle$. For

example, taking this approach the internal energy is given by

$$\overline{U} = \sum_p \varepsilon_p \big(\langle n_p \rangle - {}^1\!/_2 \big) \,. \tag{6.10}$$

The $\langle n_p \rangle$ are computed in the same iterative way we compute $\mathcal{Z}$. The procedure is mostly identical, except the p-th step of the iteration is replaced by

$$\langle n_p \rangle_p^{\mathrm{e}} = \langle n_p \rangle_{p-1}^{\mathrm{e}} \times 0 + \langle n_p \rangle_{p-1}^{\mathrm{o}} \, e^{-\beta \varepsilon_p / 2} \,, \tag{6.11}$$

$$\langle n_p \rangle_p^{\mathrm{o}} = \langle n_p \rangle_{p-1}^{\mathrm{o}} \times 0 + \langle n_p \rangle_{p-1}^{\mathrm{e}} \, e^{-\beta \varepsilon_p / 2} \,. \tag{6.12}$$

These edits to the iterative scheme, eq. (6.11) and (6.12), compute the probability of quasiparticle mode p being occupied by tracing over the other modes while respecting the parity requirement. However, they are summing the unnormalised probabilities. The final values $\langle n_p \rangle^{\mathrm{e}}$ or $\langle n_p \rangle^{\mathrm{o}}$ must be divided by their respective partition function. Since the quasiparticles are fermions the hole occupancy is automatically obtained, $\langle f_p f_p^\dagger \rangle = 1 - \langle n_p \rangle$.

Some thermodynamic quantities require further calculations. In particular, the heat capacity is related to the variance of energy, $\overline{C} = \beta^2 \mathrm{Var}(\overline{H})$. Thus, when expanded in terms of expectation values of n_p operators, it contains second order terms

$$\overline{C} = \beta^2 \sum_{p,\,q} \varepsilon_p \, \varepsilon_q \big(\langle n_p n_q \rangle - \langle n_p \rangle \langle n_q \rangle \big) \,.$$

Two-mode expectation values $\langle n_p n_q \rangle$ can again be calculated using an iterative scheme. The procedure is identical to the calculation of $\langle n_p \rangle$ apart from step q. At step q the edit applied to step p, equations (6.11) and (6.12), is repeated. This gives zero weight to the subspace where either mode p or q is unoccupied.

6.1.3. Equivalent cases

In some cases the projection does not change the thermal physics of the fermions. This section shows that if the spectrum contains a zero energy mode (which is the case for a system with an open boundary) many of the projected thermal properties are related in simple ways to their unprojected values. Most quantities of interest can then be computed

using the unprojected expressions given in section 6.1.1.

The projected partition function we have written down, eq. (6.5), enforces the parity constraint with binary function $\Pi(\underline{n})$, which kills terms in the sum that do not have the correct parity. Generalising eq. (6.5) by absorbing $\Pi(\underline{n})$ into the exponential and introducing a new Lagrange multiplier γ allows us to tune between the projected and unprojected cases,

$$\overline{\mathcal{Z}}(\beta,\gamma) = \sum_{\{\underline{n}\}} \exp\big(-\beta\,\mathcal{E}(\underline{n}) - \gamma W(\underline{n})\big) \tag{6.13}$$

where $W(\underline{n})$ is defined to be the parity of the bit-string $\underline{n}$, explicitly $W(\underline{n}) = (\sum_p n_p) \bmod 2$. (Here we have assumed for simplicity that the required parity is even. We could switch to the case where the physical states have odd parity by replacing $W(\underline{n})$ with its logical negation $\texttt{not}\, W(\underline{n})$.) In the limit $\gamma \to \infty$ this imposes a hard constraint on the parity of allowed states by assigning zero probability to states with odd parity. In the opposite limit, $\gamma \to 0$, we recover the unprojected partition function.

The two-parameter partition function, eq. (6.13), can be rewritten by splitting the set of states into its fixed parity subsets, $\{\underline{n}\} = \Pi_e \cup \Pi_o$ where $W(\underline{n}) = 0$ for all $\underline{n} \in \Pi_e$ and $W(\underline{n}) = 1$ for $\underline{n} \in \Pi_o$. This gives an expression

$$\overline{\mathcal{Z}}(\beta,\gamma) \;=\; \sum_{\underline{n}\,\in\,\Pi_e} e^{-\beta\,\mathcal{E}(\underline{n})} + e^{-\gamma} \sum_{\underline{n}\,\in\,\Pi_o} e^{-\beta\,\mathcal{E}(\underline{n})} \;\equiv\; \mathcal{Z}^e(\beta) + e^{-\gamma}\,\mathcal{Z}^o(\beta)$$

where $\mathcal{Z}^\alpha(\beta)$ is the restriction of the partition sum onto states in Π_α. We want to compare the projected and unprojected partition functions and so we divide $\mathcal{Z}(\beta,\gamma)$ evaluated in one limit by its value in the other limit,

$$\frac{\overline{\mathcal{Z}}(\beta,\gamma\to\infty)}{\overline{\mathcal{Z}}(\beta,\eta\to 0)} \;=\; \frac{\mathcal{Z}^e(\beta) + e^{-\gamma}\,\mathcal{Z}^o(\beta)}{\mathcal{Z}^e(\beta) + e^{-\eta}\,\mathcal{Z}^o(\beta)} \;\to\; \frac{\mathcal{Z}^e(\beta)}{\mathcal{Z}^o(\beta) + \mathcal{Z}^e(\beta)}\,. \tag{6.14}$$

If the spectrum contains at least one zero-mode, $\varepsilon_0 = 0$, then any state $\underline{n} \in \Pi_e$ can be mapped uniquely onto $\underline{n}' \in \Pi_o$ without changing the value of $\mathcal{E}(\underline{n})$, by flipping the occupancy of ε_0. The Boltzmann weights are unchanged by this and so $\mathcal{Z}^e(\beta) = \mathcal{Z}^o(\beta)$. Substituting this equality into eq. (6.14) we find

$$\overline{\mathcal{Z}}(\beta,\infty) = \frac{1}{2}\,\overline{\mathcal{Z}}(\beta,0)\,, \tag{6.15}$$

meaning the projected partition function can be obtained directly from the unprojected partition function. Since the thermodynamics follows from $\overline{\mathcal{Z}}$ we might expect similar relations to hold for thermal quantities. In fact, the relationship is even simpler.

In analogy to eq. (6.13) the two-parameter expectation value of an observable $\overline{\mathcal{O}}$ is

$$\begin{aligned}\langle\overline{\mathcal{O}}(\beta,\gamma)\rangle &= \overline{\mathcal{Z}}(\beta,\gamma)^{-1}\sum_{\{\underline{n}\}}\overline{\mathcal{O}}(\underline{n})\,e^{-\beta\,\mathcal{E}(\underline{n})-\gamma W(\underline{n})} \\ &= \overline{\mathcal{Z}}(\beta,\gamma)^{-1}\Bigg(\sum_{\underline{n}\,\in\,\Pi_e}\overline{\mathcal{O}}(\underline{n})\,e^{-\beta\,\mathcal{E}(\underline{n})}+e^{-\gamma}\sum_{\underline{n}\,\in\,\Pi_o}\overline{\mathcal{O}}(\underline{n})\,e^{-\beta\,\mathcal{E}(\underline{n})}\Bigg) \\ &\equiv \overline{\mathcal{Z}}(\beta,\gamma)^{-1}\big(\,\langle\overline{\mathcal{O}}(\beta)\rangle^{e}+e^{-\gamma}\langle\overline{\mathcal{O}}(\beta)\rangle^{o}\,\big)\,.\end{aligned}$$

If $\overline{\mathcal{O}}(\underline{n})$ does not change when the occupancy of the zero mode ε_0 is flipped, then we can repeat the trick used with the partition function and argue that $\langle\overline{\mathcal{O}}\rangle^{e}=\langle\overline{\mathcal{O}}\rangle^{o}$. This is the case for quantities that are functions of the energies raised to some power, such as internal energy and specific heat. Comparing the two limits we then find

$$\langle\overline{\mathcal{O}}(\beta,\infty)\rangle=\langle\overline{\mathcal{O}}(\beta,0)\rangle\,. \tag{6.16}$$

The entropy cannot be obtained this way. However we can easily reconstruct it using free energy and internal energy. Most quasiparticle occupancies will obey this condition and follow eq. (6.16). An exception is the occupancy of the zero energy mode itself $\langle n_0\rangle$, which if needed must be computed iteratively.

6.2. Aside – Ground state topological degeneracy

Both the toric code and Ising anyon models have topologically degenerate ground states. These degeneracies are lifted by topologically non-trivial loops of operators, so when defined on a torus there are four *topological sectors*. For the toric code all four of these are degenerate (as discussed in chapter 2). In contrast, the topological quantum field theory (TQFT) description of the non-Abelian Ising anyons has only a three-fold degenerate ground state. This section shows how these topological degeneracies manifest in the fermionic physics of the no-vortex sector.

Homologically distinct gauge field configurations (corresponding to the different topolog-

ical sectors) can be generated by flipping link variables along non-trivial loops. Specifically, the link variables are flipped $u_{ij} \to -u_{ij}$ on z bonds along a horizontal row or on x bonds down a vertical column as illustrated in fig. 6.1 (a) and (b). These configurations cannot be generated by local gauge transforms D_i and are equivalent to the gauge configurations that would be produced by creating a pair of vortices, transporting one of them around either of the torus handles and annihilating the pair. Combinations of these loops generate all four homologically distinct configurations.

Inserting non-trivial gauge loops into the no-vortex sector generates four distinct fermionic Hamiltonians. The eigenenergies of each should be identical, but in practice this does not hold precisely true. There are small differences between the fermionic spectra of the different topological sectors. However, these differences are seen to vanish as $L \to \infty$. To make this precise, define a normalised vector $\hat{\varepsilon} = (\varepsilon_1, \varepsilon_2, \ldots)/\sqrt{\varepsilon_1^2 + \varepsilon_2^2 + \ldots}$ from the positive fermionic eigenenergies ε_p. Figure 6.1(c) plots the (Euclidean) distance between the vectors defined from the no-loop spectrum $\hat{\varepsilon}_0$ and that produced by introducing an x-loop $\hat{\varepsilon}_x$, as a function of system size. The data displayed is for the Ising phase, with similar results being found for the toric code phase. The distance is seen to decay as a power law $\|\hat{\varepsilon}_0 - \hat{\varepsilon}_x\| \sim L^{-a}$. Interestingly, the differences are primarily in the structure of the spectrum. The difference between the fermionic ground state energies $E_G = -\sum_p \varepsilon_p/2$ vanishes exponentially with L as shown in the inset of fig. 6.1(c).

Finite size effects aside, the main determinant of the ground state degeneracy is the fermion parity of each topological sector. Since the no-vortex sector spectrum is gapped odd parity states (which must contain an excitation) have higher energy, meaning the topological degeneracy is indicated by the number of sectors that have even physical parity. In the toric code phases the physical parity is found to be even for all four giving a four fold topological degeneracy. Whereas, in the Ising phase three of the four have even fermion parity but the fourth has odd fermion parity. These differences are highlighted by plotting the projected thermal fermion energies as $T \to 0$ in fig. 6.2. As $L \to \infty$ the finite size effects vanish and these topological degeneracies become exact, matching the expectations of TQFT.

This calculation requires careful handling of the diagonalisation, as discussed in note B.1.2. Additionally, it is only seen to hold for even system sizes. For square systems of size $L \times L$

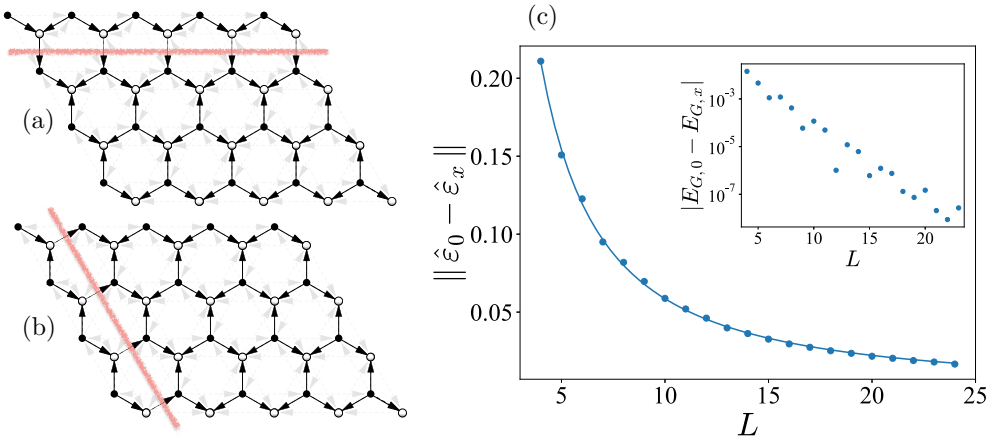


Figure 6.1. – Non-trivial loops in the (a) x and (b) y lattice directions are generated by flipping all of the link variables along a row or column. The arrows are redrawn to show the new gauge configurations and coloured lines are added to indicate where the flips in the directions of the arrows have occurred. (Periodic boundary conditions are assumed, but the links completing the periodic boundaries are not drawn.) (c) the spectral distance (defined in the text) between the no-loop and x-loop fermionic spectra in the B phase as a function of system size L, fit to a power law decay $\|\hat{\varepsilon}_0 - \hat{\varepsilon}_x\| = 1.46\, L^{-1.40}$. (inset) the difference between the ground state energies of the spectra decays exponentially. Data is displayed for parameters $J_x = J_y = J_z = 1$ and $K = 0.1$.

where L is odd, we find the toric code phase has a two-fold degenerate ground state. The Ising phase for odd L flips the picture we previously described, with three of the four configurations becoming odd and a single even, giving a unique ground state. A similar even/odd effect is seen in the degeneracy of the Moore-Read state. This wavefunction is believed to describe fractional quantum Hall systems in a certain regime, and also supports Ising anyon quasiparticles [40]. The Moore-Read state of an odd number of fermions has unique degeneracy, whereas for an even number of fermions it has three-fold degeneracy [100].

To the best of my knowledge, this relationship between topological degeneracy and fermion parity has not previously been demonstrated for the Kitaev fermionisation. Similar results have been found in another fermionisation of the honeycomb [101–103]. There, through an alternative Jordan-Wigner like transformation, the fermionic part of the honeycomb ground state is expressed as a BCS state. The BCS state is an even fermion number parity state. The different topological sectors are investigated by imposing periodic and anti-periodic boundary conditions on the BCS wavefunction and the authors find that in the Ising phase one of the four BCS states vanishes. Thus giving the ground

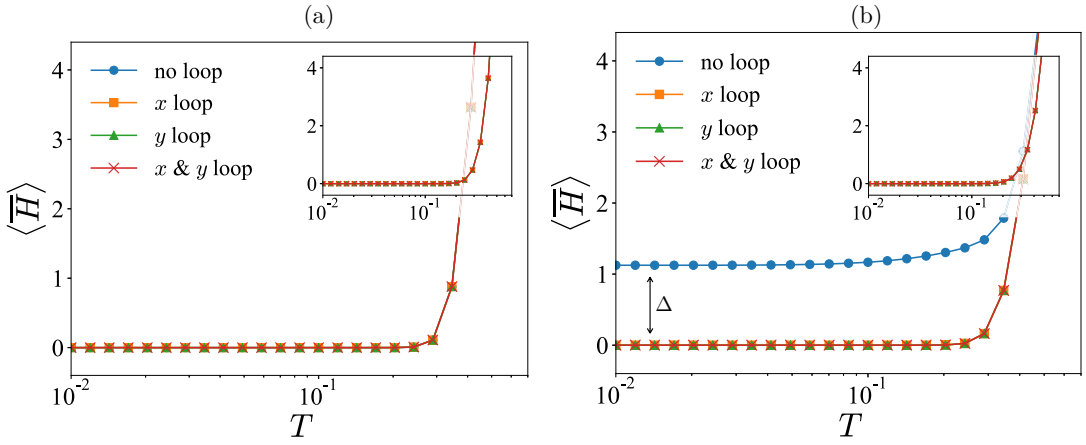


Figure 6.2. – Projected thermal fermion energy for the four homologically distinct gauge field configuration of the no-vortex sector in the (a) A phase, and (b) B phase. (Inset) the thermal fermion energies computed without projection. (a) all four homology classes are degenerate. (b) as $T \to 0$ one of the four has odd fermion parity and so tends to a higher energy, with the splitting equal to the fermion gap Δ. Computed for a 16×16 system with parameter values $K = 0.1$ and (a) $J_x = J_y = 0.1$, $J_z = 1$, (b) $J_x = J_y = J_z = 1$.

state a three-fold topological degeneracy. In contrast, in the toric code phase they all exist giving a four-fold degeneracy.

6.3. Thermal distributions of vortices

The fermionisation of the spin model separates out the physical degrees of freedom of the model into fermions, whose thermal physics was addressed in the last section, and vortices. While we were able to solve exactly for the thermal physics of the fermions, it is not possible to do the same for the vortices and we must use approximate methods.

The vortex data is encoded in the gauge fields u_{ij}, which are captured by matrix A_{ij} of the fermionic Hamiltonian $\overline{H}$. We will now make the dependence explicit by writing $\overline{H}(u)$. The partition function of the full model is a trace over all degrees of freedom. This can be expressed as a sum over all gauge configurations of the fermionic partition functions,

$$\mathcal{Z}(\beta) = \sum_u \mathrm{tr}\left(\mathcal{P} e^{-\beta \overline{H}(u)}\right) \equiv \sum_u \overline{\mathcal{Z}}(\beta; u)\,. \tag{6.17}$$

Here $\overline{\mathcal{Z}}(\beta; u)$ is the partition function of the fermions, studied in the previous section. We numerically study the thermodynamics described by $\mathcal{Z}$ using Monte Carlo methods. The

u fields are gauge degrees of freedom and the fermionic spectrum should only depend on the locations of the vortices. This makes it tempting to replace the sum over u with a sum over vortex patterns. However, as was shown in section 6.2, non-trivially different gauge configurations can have different physical fermion parities and therefore contribute differently to eq. (6.17). This means we must be careful when restricting the sum.

6.3.1. Monte Carlo methods

Many quantities of interest correspond to the expectation value of some observable of the full model $\mathcal{O}$. These observables might depend only on the vortices, for example the vortex density, or they may additionally depend on the fermionic state, such as energy. Generically these expectation values can be written

$$\langle \mathcal{O}(\beta) \rangle \;=\; \sum_{u} \frac{\mathcal{Z}(\beta;u)}{\boldsymbol{\mathcal{Z}}(\beta)} \, \langle \overline{\mathcal{O}}(\beta;u) \rangle \,, \tag{6.18}$$

where, when applicable, $\langle \overline{\mathcal{O}}(\beta;u) \rangle$ means the fermionic expectation value of the observable for gauge configuration u, evaluated using the expressions or methods described in the previous section. Sums like eq. (6.18) can be approximated by choosing a finite set of M gauge configurations $\{u_1, u_2, \ldots u_M\}$ where each u_i occurs with the correct probability, which is $\overline{\mathcal{Z}}(\beta;u_i)/\mathcal{Z}(\beta)$, and computing a mean,

$$\langle \mathcal{O}(\beta) \rangle \;\approx\; \frac{1}{M} \sum_{i=1}^{M} \langle \overline{\mathcal{O}}(\beta;u_i) \rangle \,. \tag{6.19}$$

We select this set of M gauge configurations using a Markov chain approach, discussed in detail below. This covers many quantities, however, not all thermodynamic properties can be computed as the expectation value of an observable.

In particular, the heat capacity of the full model is not the average of the fermionic specific heats. As always it is given by $C = \beta^2 \mathrm{Var}(\, H \,)$, where H is the full Hamiltonian. The fermionic and vortex degrees of freedom can be separated out and C written

$$C(T) \;=\; \big\langle\, \overline{C}(u) \,\big\rangle \;+\; \frac{1}{T^2} \mathrm{Var}\big[\, \overline{U}(u) \,\big] \,, \tag{6.20}$$

where $\overline{C}(u)$ and $\overline{U}(u)$ are the fermionic specific heat and internal energy for gauge configuration u. The mean $\langle \cdot \rangle$ and variance $\mathrm{Var}[\cdot]$ are computed with respect to the distribution over gauge configurations. The variance of $\overline{U}(u)$ can be estimated from the same sample $\{u_1, u_2, \ldots u_M\}$. Entropy is computed from the specific heat as

$$S(T) = N \ln(2) - \int_T^\infty \mathrm{d}T' \, \frac{C(T')}{T'} \,, \tag{6.21}$$

which follows from the infinite temperature limit of the model being N independent spin-1/2 particles. We next describe how we select the u in the sample.

Markov chain

The aim is to choose M gauge configurations such that the probability of any u_i being chosen is $\overline{\mathcal{Z}}(\beta; u_i)/\mathcal{Z}(\beta)$. We achieve this with a Markov chain approach. Beginning with the uniform $u_{i,j} = +1$ gauge configuration (which is in the no-vortex sector), at each step a new gauge configuration u' is proposed. The proposal will be described in detail below. Each time the current u is replaced by the proposed u' with probability

$$\mathrm{Prob}(\text{ accept } u \to u') \;=\; \min\left[1, \, \frac{\overline{\mathcal{Z}}(\beta; u')}{\mathcal{Z}(\beta)} \Big/ \frac{\overline{\mathcal{Z}}(\beta; u)}{\mathcal{Z}(\beta)} = \frac{\overline{\mathcal{Z}}(\beta; u')}{\overline{\mathcal{Z}}(\beta; u)} \right] . \tag{6.22}$$

After a large number of attempted replacements ($\sim L^2$) we add the current gauge configuration to the sample, this timescale we refer to as a *sweep*. This continues until we have collected enough configurations. Each Markov step is relatively slow as computing the acceptance probability requires rediagonalising the fermionic Hamiltonian. We additionally use parallel tempering to help speed up the convergence of statistical estimates [104].

The proposal $u \to u'$ is generated by selecting two random plaquettes and drawing a string of gauge flips between them, fig. 6.3(a). This might select the same plaquette twice, in which case the string is empty. The string is combined with a randomly selected topologically non-trivial loop of gauge flips fig. 6.3(a)-(d). This choice of proposal effectively ignores the trivial gauge loops within eq. (6.17) and is sampling only over vortex sectors and non-trivially different gauge configurations.

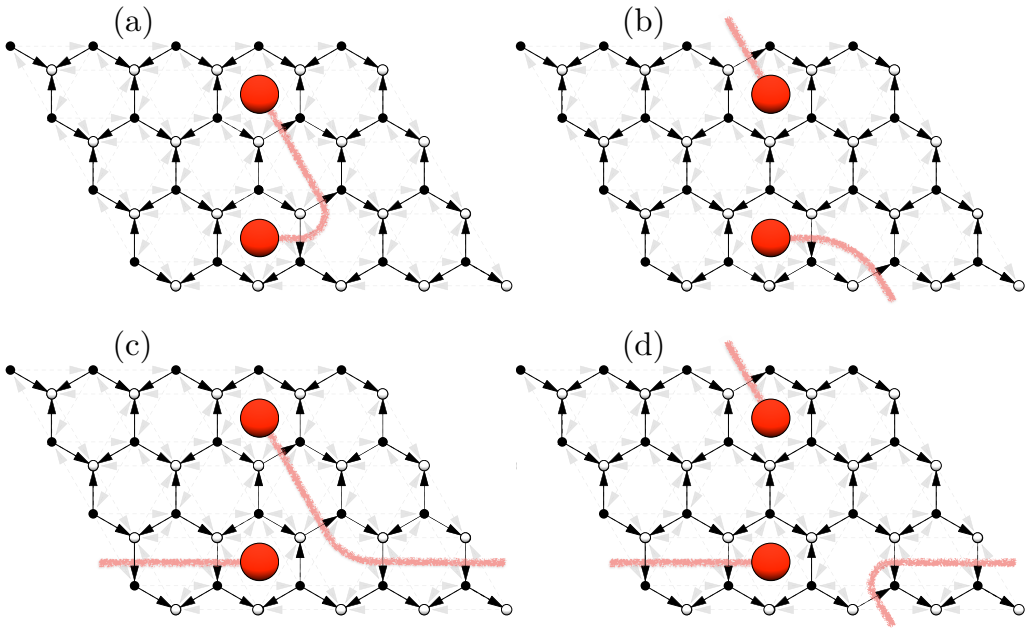


Figure 6.3. – Markov step proposal. (a) two plaquettes are selected at random and an arbitrary rectangular path of flipped links is drawn between them. With equal probabilities, this gauge string is combined with (a) nothing, (b) a non-trivial loop wrapping the lattice up-down, (c) a non-trivial loop wrapping the lattice left-right, or (d) both of these non-trivial loops.

Parallel tempering

Using parallel tempering has two advantages. It both speeds up the convergence of statistical estimates such as eq. (6.19) and additionally it computes these estimates for a set of different temperatures in a single simulation run. Multiple Markov chains, called *replicas*, are evolved simultaneously at R different temperatures $\{\beta_1, \beta_2, \ldots \beta_R\}$. After each replica has undergone some number of sweeps the replicas have a probability of being exchanged. Exchanges happen on a nearest neighbour basis. So as to allow full mixing, we randomly choose between (β_1, β_2), $(\beta_3, \beta_4), \ldots$ or (β_2, β_3), $(\beta_4, \beta_5), \ldots$ pairings to initiate swaps. Replicas are exchanged within each pair with probability

$$\text{Prob}(\, \beta_i \leftrightarrow \beta_{i+1} \,) \;=\; \min \left[1, \frac{\overline{\mathcal{Z}}(\beta_{i+1}; u^i)\; \overline{\mathcal{Z}}(\beta_i; u^{i+1})}{\overline{\mathcal{Z}}(\beta_i; u^i)\; \overline{\mathcal{Z}}(\beta_{i+1}; u^{i+1})} \right] , \tag{6.23}$$

where u^i is the gauge configuration of the replica currently at temperature β_i. This allows the replicas to diffuse up and down in temperature and the Markov chains to more efficiently explore the state space. The sample set of gauge configurations $\{u_1^i, u_2^i, \ldots u_M^i\}$ at fixed temperature β_i is compiled between swaps. For example, if swaps occur after

every sweep, the current vortex configuration of replica i is added to the sample before deciding swaps. Further, parallel tempering can be implemented without any real increase in runtime by making use of distributed parallel computing (discussed in note B.1.4).

6.4. Finite temperature properties

In this section we use the methods described in sections 6.1 and 6.3 to study the thermodynamics of the spin model. Monte Carlo simulation data is collected for $\sim 50,000$ samples, employing varying numbers of parallel tempering replicas. The initial portion of data prior to equilibration of the simulations is discarded then averages, variances and error estimates are computed from the remaining series. The analysis performed on the data, including obtaining these estimates and deciding equilibration time, is described in detail in in appendix B.2. Toric code plots correspond to parameters in the A phase of the no-vortex sector, and what is referred to as the Ising regime corresponds to parameter values in the B phase of the no-vortex sector. In this section we only study even system sizes. This simplifies the discussion by allowing us to avoid the even/odd effects highlighted in section 6.2.

We begin by studying the thermodynamic properties of the system such as the internal energy, specific heat and thermal entropy in section 6.4.1. Many of these quantities have been studied extensively in previous published works [60, 105]. Our results agree with the findings of these studies where they overlap. The main new contribution we make in this section is to carefully compute the thermal entropy down to low temperatures. We show that as $T \to 0$ the thermal entropy yields the topological ground state degeneracies of the different regimes of the model. It is not possible to show this without fully accounting for the fermionic parity, which was not done in previous studies. We also highlight interesting differences in the behaviour of the fermionic parity at higher temperatures, which relates to the different characters of the vortices in the toric code and Ising regimes. Following this, in section 6.4.2 we study the fermionic gap as a function of temperature. This provides a intuitive explanations for some of the differences between the thermal behaviours of the two regimes. Finally, in section 6.4.3 we address one of the most striking differences between the two regimes of the model: the emergence of a thermal metal phase in the

Ising regime at high temperature. By studying the distribution of fermionic energy levels and the spatial properties of the wavefunctions we provide clear evidence for the existence of this phase.

6.4.1. Thermodynamics

To analyse the thermal behaviour of the spin model we calculate the specific heat, *i.e.* the heat capacity per particle, C/N, thermal entropy, S, vortex density, ρ_V, internal energy per particle, U/N ($U = \langle H \rangle$), and average fermion parity, $\langle \mathcal{F} \rangle$. These quantities are plotted against temperature for both the toric code and Ising coupling regimes in fig. 6.4.

We first notice the distinctive two peak structure of the specific heat, indicating the physical fractionalisation of the spins, that has been observed in previous studies [60]. For both the toric code and Ising regime, the low T peak is associated with a change of the vortex density, in which it rises from its ground state value $\rho_V = 0$ towards a completely disordered vortex configuration $\rho_V = 1/2$. The higher temperature peak is linked to a large change in internal energy U, corresponding to a proliferation of fermionic quasiparticles. These peaks are well separated by several orders of magnitude of T in the toric code limit, whereas they are closer and blur together somewhat in the Ising regime. Nevertheless, the specific heat in both regimes exhibits qualitatively similar behaviour.

In the upper panel of fig. 6.4 we plot the behaviour of the thermal entropy, S. We see that at small temperatures the thermal entropy identifies the topological degeneracy of the ground state. Specifically, the Boltzmann formula gives

$$S = \ln(\Omega) \tag{6.24}$$

where Ω is the number of microstates (with the Boltzmann constant set equal to one). As $T \to 0$, Ω gives the degeneracy of the ground state. Hence, the thermal entropy we compute is in excellent agreement with the topological degeneracies expected for the toric code anyons $S = \ln 4 = 2 \ln 2$ and for the Ising anyons $S = \ln 3$. However, for larger temperatures when the vortices start to proliferate these entropic differences between the different coupling regimes are washed out.

In contrast to the thermal entropy, the fermionic parity $\langle \mathcal{F} \rangle$ shows interesting differences

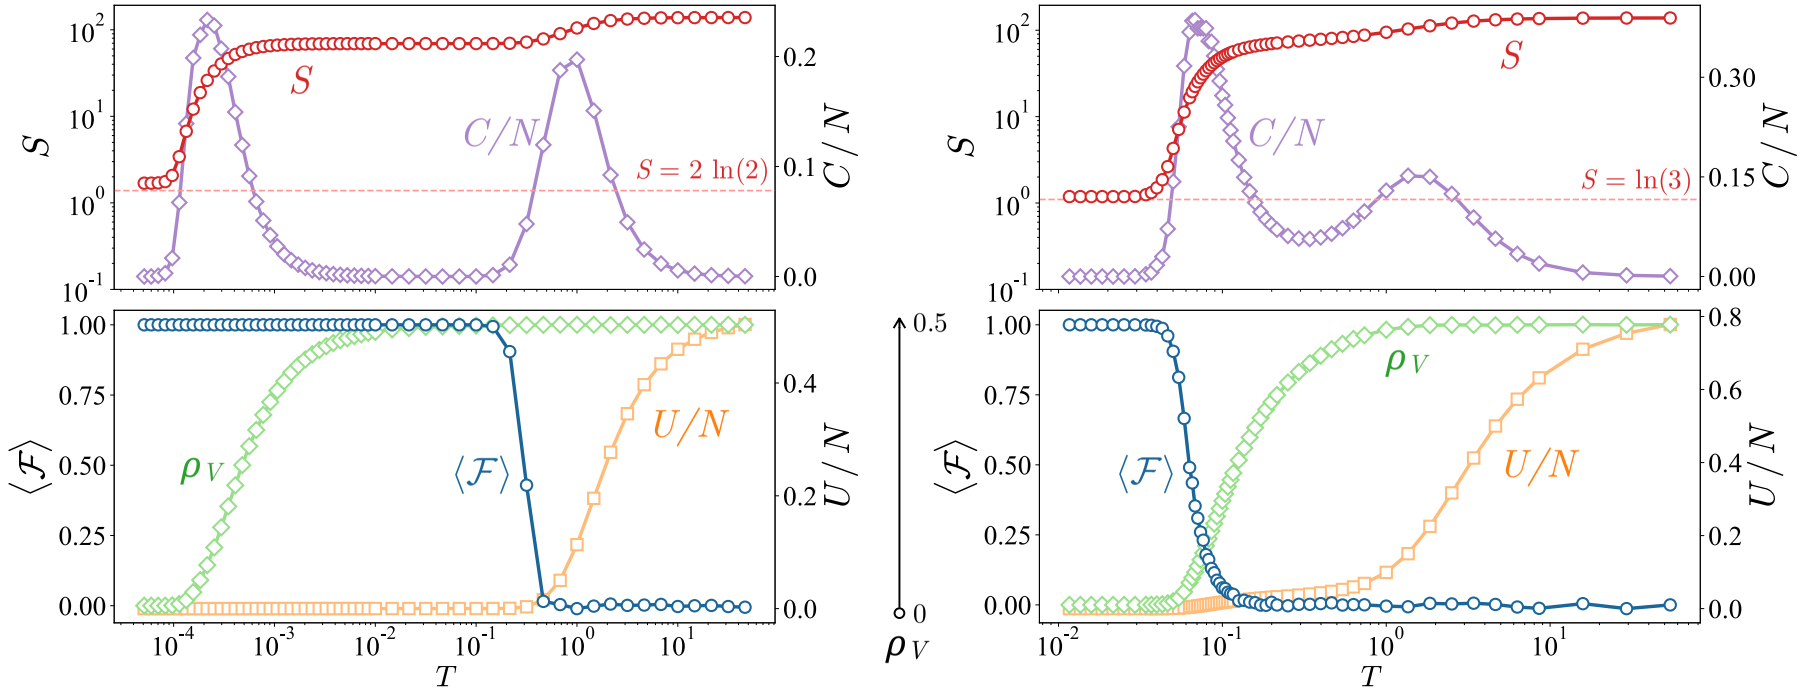


Figure 6.4. – Thermodynamic quantities in the (left) toric code and (right) Ising regime. In the top panel we plot the entropy S (red circles) and the specific heat C/N (purple rhombus) of the model. The distinctive two peak structure of C/N is clearly visible, with the lower temperature peak associated with the vortex excitations and the higher temperature peak with the fermionic excitations. At low temperatures S identifies the different topological degeneracies of the ground states, which is fourfold in the toric code case ($S = 2\ln 2$) and threefold in the Ising case ($S = \ln 3$). At higher temperatures the entropic signature of topological order is lost due to the proliferation of vortices. The lower panel plots vortex density ρ_V (green rhombus), internal energy per particle U/N (orange squares) and average fermion parity $\langle\mathcal{F}\rangle$ (blue circles). The vortex density ρ_V is zero at low temperature, and increases to $1/2$ when the first peak of C/N occurs, while the internal energy becomes non-zero at the second peak. The average fermion parity $\langle\mathcal{F}\rangle$ shows distinctly different behaviours in the Toric code and Ising regimes for intermediate temperatures between the two peaks. Data is shown for $L = 10$ system size, with couplings $J_x = J_z = 0.25$, $J_z = 1$, $K = 0$ in the Abelian case and $J_x = J_y = J_z = 1$, $K = 0.1$ in the non-Abelian case.

between the two regimes at higher temperatures and in fact hints at even more complex differences between the two. We begin by describing the behaviour we would expect $\langle \mathcal{F} \rangle$ to have based purely on the anyon models. The $T \to 0$ state always has no vortices, in the toric code case this already rules out the possibility of odd fermion parity states because of the fusion rules. (States containing two fermionic quasiparticles would be allowed, however these have very high energy.) Whereas in the Ising regime odd fermion parity states that do not contain any vortices are allowed by the fusion rules (they correspond to states with anyonic flux around both handles of the torus as studied in section 6.2). However, since they contain a fermion they have higher energy and so are exponentially suppressed in the thermal state at low temperature. Hence at low T we expect to find $\langle \mathcal{F} \rangle = 1$ in both regimes. In contrast, at high temperatures we would expect different values of $\langle \mathcal{F} \rangle$ in the two cases. As $T \to \infty$ the vortices becomes completely disordered, $\rho_V = 1/2$. In the toric code case half of these configurations will contain odd numbers of e (and m) anyons and half even numbers. From the fusion rules these correspond to odd and even fermion parities respectively, and averaging over them gives $\langle \mathcal{F} \rangle = 0$. Whereas in the Ising theory the only way to generate a state with total odd fermion parity is to add non-trivial anyon flux around both torus handles. This is one of four flux configurations and so we would instead expect to find $\langle \mathcal{F} \rangle = {}^3\!/_4(+1) + {}^1\!/_4(-1) = 0.5$. Figure 6.4 (lower panel) plots the fermion parity as a function of temperature. We see that at low temperature both regimes agree with the predictions of the anyon models. At high temperature, however, only the toric code case agrees with our predictions. The fermion parity also becomes zero in the Ising case at high temperatures. This is a first indication that the thermal proliferation of Ising vortices has driven a transition out of the topologically ordered phase (described by the anyon model) to a new kind of phase. We will demonstrate that this phase is a thermal metal.

It is also interesting that the temperatures at which $\langle \mathcal{F} \rangle \to 0$ are different in the two regimes. In the Ising case it corresponds with the change in the vortex density (proliferation of vortices) and in the toric code case with the change in the internal energy (proliferation of fermions). We can explain this difference by considering the behaviour of the fermion gap.

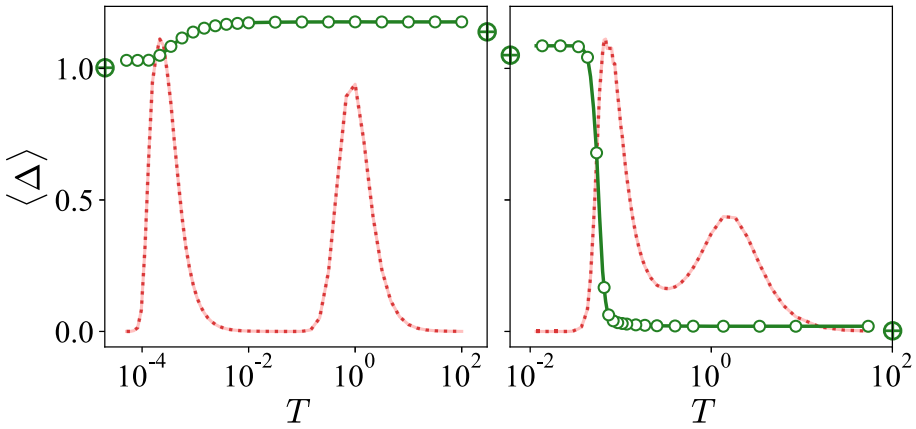


Figure 6.5. – Average fermion gap for the (left) toric code and (right) Ising cases, with the specific heat profile superimposed. The fermions in the toric code limit are gapped at all temperatures, whereas in the Ising case they become gapless when vortices start proliferating. The large system size limits of the fermion gap for $T \to 0$ and $T \to \infty$ are indicated with markers on the axes. Data is shown for $L = 10$ system size with couplings set to $J_x = J_z = 0.25$, $J_z = 1$, $K = 0$ in the toric code case and $J_x = J_y = J_z = 1$, $K = 0.1$ in the Ising case.

6.4.2. Fermionic gap

The average fermion gap $\langle \Delta \rangle$, defined as the lowest lying fermionic energy averaged over the vortex sectors that are thermally excited for a given temperature T, helps to explain some of the differences we observe between the two regimes. Additionally, we can use it as a useful check on our Monte Carlo simulation data. It is simple to sample from the zero temperature and infinite temperature gauge configurations. Here we exploit this to test our Monte Carlo results by comparing these limiting cases to the temperature extremes of our simulation data.

We plot $\langle \Delta \rangle$ against temperature for the toric code and Ising regimes in fig. 6.5. We see that the toric code regime remains gapped at all temperatures, whereas in the Ising case the average fermionic gap vanishes above the temperature at which vortices begin to proliferate. This explains some of the qualitative differences we observe between the behaviour of the fermion parity in the two regimes, fig. 6.4. Since the fermions of the toric code case remain gapped at all temperatures, vortex sectors that correspond to odd fermion parity are energetically suppressed until the second specific heat peak is reached. In contrast, since the fermions of the Ising regime become gapless at the first specific heat peak, above that temperature there is no energetic penalty to odd-parity vortex sectors and the fermion parity can take on new values $\langle \mathcal{F} \rangle < 1$ immediately.

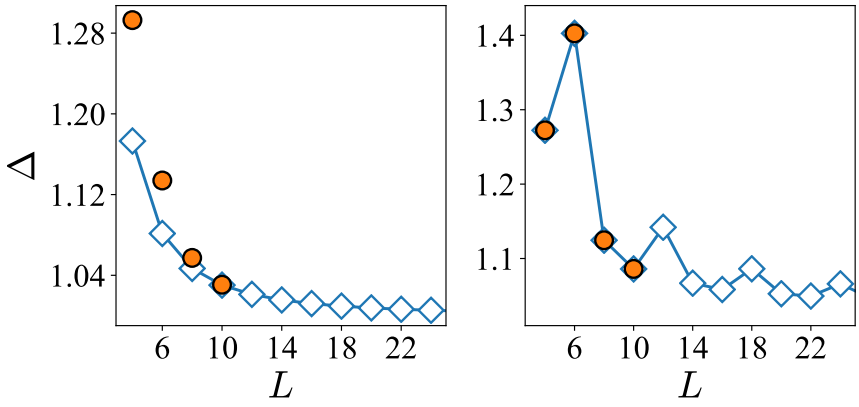


Figure 6.6. – Zero temperauture fermion gap as a function of system size for (left) toric code, $J_x = J_z = 0.25$, $J_z = 1$, $K = 0$, and (right) Ising, $J_x = J_y = J_z = 1$, $K = 0.1$, couplings. Both cases quickly decay towards a fixed value. The fermion gaps obtained from the lowest temperature Monte Carlo simulations are displaying (orange circles) on the plots.

Zero temperature

The ground state of the model lives in the no-vortex sector [51]. This ground state has a topological degeneracy on a torus, which corresponds to the four different choices of periodic/anti-periodic boundary conditions for the fermions [102]. Due to the projection procedure, however, some of these topological sectors may have odd fermion parity and so higher ground state energy. We take the zero temperature gauge configuration to be an equal weight mixture of the even parity topological sectors, i.e. Δ is computed as the mean of the fermion gaps of the even parity no-vortex sectors. In the Ising case this is three of the four sectors, whereas for the toric code all four are degenerate. We will ignore the finite size gapping of the topologically degenerate states, as this vanishes exponentially.

The fermion gap of this zero temperature gauge configuration is plotted as a function of system size for the Ising and toric code cases in fig. 6.6. Both are decaying to some fixed value, and the Ising case additionally displays oscillations. Extrapolating these trends we find the limiting value of the zero temperature fermion gap.

The lowest temperature value of the gap obtained from our Monte Carlo simulation is plotted alongside the curves in fig. 6.6. We find good agreement between the Ising Monte Carlo data and the zero temperature gauge configuration. However, our toric code data breaks away from it for small systems, $L < 10$. This occurs as a result of the small vortex gap in the toric code case, recall that the first specific heat peak is at $T \sim 10^{-4}$. Here $1/T$

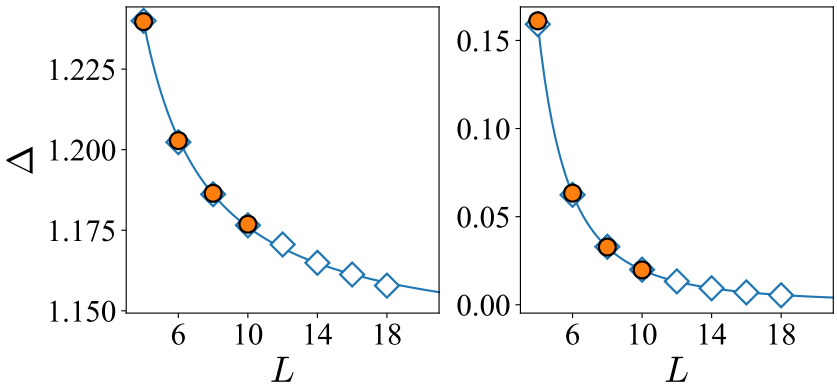


Figure 6.7. – Infinite temperauture fermion gap as a function of system size for (left) toric code, $J_x = J_z = 0.25$, $J_z = 1$, $K = 0$, and (right) Ising, $J_x = J_y = J_z = 1$, $K = 0.1$, couplings. Both decay quickly towards a fixed values, which from polynomial fits are found to be $\Delta = 1.13(9)$ in the toric code case and $\Delta = 0.000(6)$ in the Ising case. The fermion gaps obtained from the highest temperature Monte Carlo simulations are displaying (orange circles) on the plots.

is comparable to the finite size gapping of the topologically degenerate ground states for $L < 10$ and so the Monte Carlo simulation resolves the degeneracy. Our assumption that the finite size gapping of the topological degeneracy can be neglected fails in these cases. This demonstrates that $L = 10$ is the smallest system size we can consider.

Infinite temperature

Large temperatures can be studied by sampling from completely random gauge configurations. fig. 6.7 shows the fermion gaps obtained by averaging over 10,000 random u patterns as a function of system size L. We fit a polynomial decay to the data. In the toric code case the trend is toward a non-zero gap, whereas in the Ising case the gap vanishes in the large system limit. The Monte Carlo values obtained at the highest temperatures studied are plotted on the graphs and we find excellent agreement.

6.4.3. Thermal metal behaviour

Thermal metals are an exotic phase of disordered two-dimensional superconductors that break both time reversal and spin rotation symmetries, such as spinless p-wave superconductors [106, 107]. This phase is truly metallic, in the sense that its excitation spectrum is gapless about zero energy. Thermal metals have many surprising properties such as a bulk thermal conductivity that diverges with system size and nonlocalised low energy

states whose wavefunctions resemble those found at the critical point in systems showing Anderson localisation [108, 109].

Thermal metals are conductors of energy (or heat) via their fermionic degrees of freedom. A vanishing fermion gap is thus a necessary condition for a thermal metal. In the previous section we verified that this was the case for the Ising regime at temperatures above the vortex activation energy. However, since the Ising vortices bind Majorana modes at zero-energy the vanishing of the gap is not sufficient. It cannot distinguish between a metallic phase (which has fermionic energy levels occuring continuously down to zero energy) and a gapped insulating phase with additional zero energy localised states associated with the vortices (isolated subgap fermionic states at zero energy). Further, a thermal metal phase is expected to display its own distinct behaviours in both the distribution of energy levels at low energy and the spatial properties of the wavefunctions [108, 109]. In the following we probe both of these properties in detail.

The fermionic distribution of energies (DOE) is given by $\rho(E) = \langle \sum_n \delta(E-\varepsilon_n) \rangle$, where the ε_n are the fermionic energies in a single vortex sector and the average is taken over the vortex sectors visited at temperature T. (Viewing the system as an ensemble of disordered fermionic Hamiltonians with quenched vortex disorder this would be the ensemble averaged density of states. This is the situation in which thermal metals are usually discussed. However, since the vortex disorder we consider arises from the thermal state of the Kitaev spin Hamiltonian, we avoid that term so as not to confuse $\rho(E)$ with the density of states of that underlying Hamiltonian.) In a thermal metal $\rho(E)$ diverges logarithmically with decreasing energy at low energies [106, 109]. To clearly see this feature of $\rho(E)$ we find it necessary to go to larger systems than are accessible when using the algorithm that properly samples from the thermal state, e.g. our $L = 10$ data. For system size $L = 20$ data is obtained by randomly sampling over vortex sectors, matching the average vortex density to the value obtained from smaller system sizes. For high temperatures, above the vortex peak of the specific heat, this has been shown to work well [110]. The logarithmic divergence is extracted by a fit to $L = 20$, and is in agreement with the behaviour of $L = 10$, as shown in fig. 6.8. For the very lowest energies $\rho(E)$ additionally exhibits characteristic oscillations predicted by random matrix theory [106, 111]. In this limit the

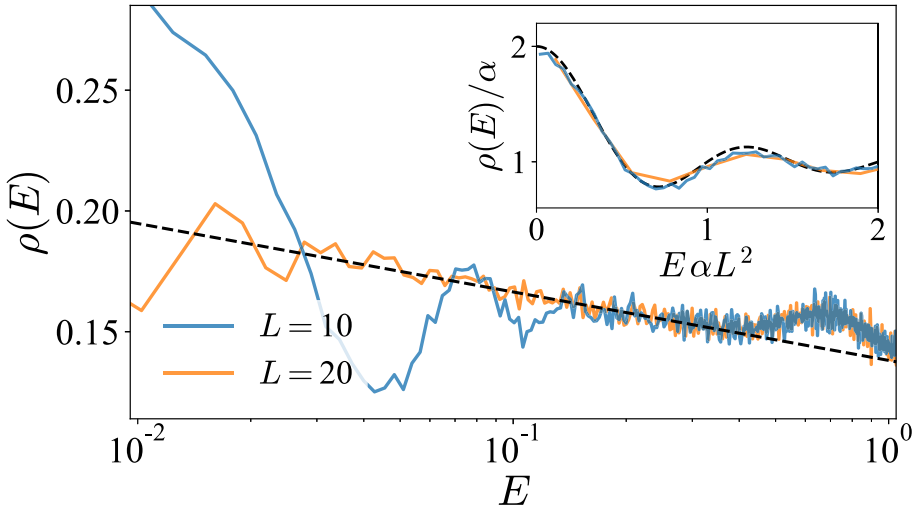


Figure 6.8. – Distribution of fermion energies $\rho(E)$ for low E in the non-Abelian case. Computed for system sizes $L = 10$ and 20 at fixed temperature $T \approx 0.4$. At low energies $\rho(E)$ diverges logarithmically with decreasing E. This is indicated by the black dashed line, which is a fit to the $L = 20$ data. For the smallest values of E oscillations are visible. (Inset) The oscillations are collapsed onto curves of the form Eqn. (6.25). The system couplings are set to $J_x = J_y = J_z = 1$, $K = 0.1$.

DOE is given by

$$\rho(E) = \alpha \left(1 + \frac{\sin(2\pi E \alpha L^2)}{2\pi E \alpha L^2}\right) . \tag{6.25}$$

The parameter α can be obtained by numerically solving the equation $\rho(1/\alpha L^2) = \alpha$ and α is expected to diverge logarithmically with L [109]. The first few oscillation periods of the random matrix theory oscillations, with α inferred from the data as just described, are shown in the inset of fig. 6.8. Hence, the energy distribution is in excellent agreement with the thermal metal phase.

Beyond the energy eigenvalues the wave functions of thermal metals exhibit *weak multifractality*, which means they are spatially extended in a particular way [112]. This behaviour is shared with systems that exhibit Anderson localisation, where weak multifractality occurs at the critical point between insulating and conducting phases [107]. Whether states are localised or extended can be studied via the inverse participation ratios of the fermionic eigenfunction $\psi_n(\mathbf{r})$ [107], given by

$$I_q = \int \mathrm{d}^2 r \, |\psi_n(\mathbf{r})|^{2q} \sim \frac{1}{L^{\tau(q)}}. \tag{6.26}$$

The exponent is often written $\tau(q) = (q-1)D_q$ where D_q is called the *fractal dimension*.

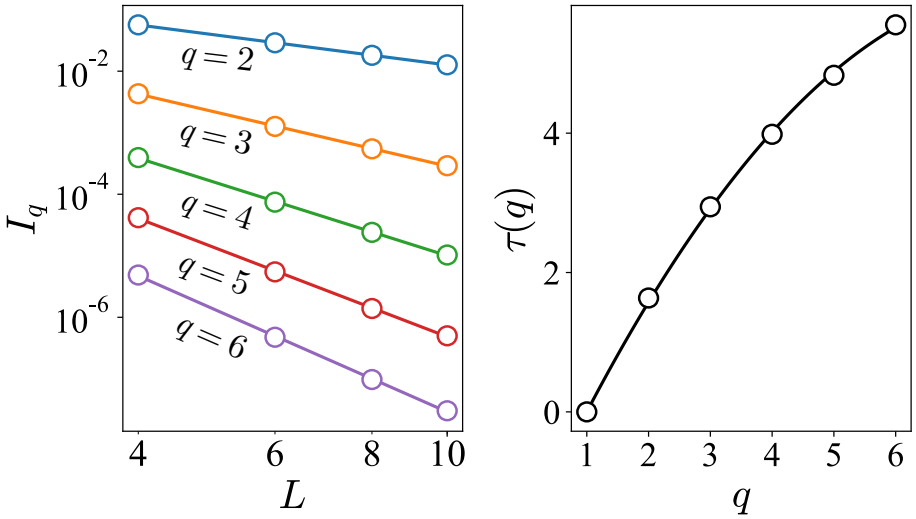


Figure 6.9. – Weak multifractality of the fermionic wavefunctions in the non-Abelian case at high temperatures. (Left) Inverse participation ratios I_q of the lowest-lying fermionic eigenmode as a function of system size L (on a log-scale). The I_q are averaged over the thermal distribution of vortices at $T \approx 0.4$. (Right) Exponents $\tau(q)$ (defined in Eq.(6.26)) as a function of q demonstrating multifractal critical behaviour, in-between the metallic $\tau(q) = 2(q-1)$ and the insulating $\tau(q) = 0$ behaviours.

For pure metallic states D_q is equal to the spatial dimension, $D_q = 2$, and for entirely localised states $D_q = 0$. However, weak multifractality implies D_q becomes dependent on q and behaves as $D_q = 2 - \gamma q$, where γ is a non-universal constant [112]. The inverse participation ratios of the lowest-lying fermionic state, obtained for temperature $T \approx 0.4$, are plotted as a function of system size for the first few values of q in the left panel of fig. 6.9. This is computed for system sizes up to $L = 10$, only using the thermal sampling of gauge configurations. The right panel of fig. 6.9 shows the exponent $\tau(q)$ extracted from fits this data, plotted as a function of q. Fitting to the data we obtain the fractal dimension $D_q = 1.7(8) - 0.12(3)q$.

6.5. Conclusions

This chapter has studied the thermal physics of the Kitaev honeycomb. We have worked from the Kitaev fermionisation and explicitly carried out the projection of the fermions into the physical subspace. The Kitaev fermionisation allows us to easily access the four homologically distinct gauge field configurations on a torus and examine the topological degeneracy of the ground state. We were able to compare the degeneracy of the Kitaev

honeycomb ground state with the degeneracies expected from the abstract description of topological quantum field theory. Up to finite size corrections, we found agreement for even system sizes. Differences for odd system sizes are consistent with previously obtained results for Ising anyons in another setting [100].

Using Monte Carlo simulations we then studied the thermal distributions of vortices. We find thermodynamic behaviours that agree with past studies of the honeycomb model at finite temperatures [59, 60]. By explicitly accounting for the total fermionic parity, which is fixed to be either even or odd depending on the gauge configuration, we can go beyond the results of these previous studies. In particular, we have shown that the topological degeneracies of the different regimes of the model can be extracted from the thermal entropy at low temperatures. We also highlighted the interesting properties of the fermion parity as a finite temperature observable. The fermion parity shows clear differences between the different regimes of the model at higher temperatures, which we explained by studying the finite temperature behaviour of the fermion gap. Following this we carried out an in depth analysis of the high temperature behaviour of the Ising regime and clearly showed that it behaves as a thermal metal. In order to do this we studied the distribution of fermionic energies and the inverse participation ratios of the low lying fermionic states.

Thermal metal phases have been seen in other similar contexts. In particular, the Kitaev honeycomb with high levels of randomness in the couplings (strong enough to introduce vortices) has previously been numerically shown to support thermal metal phases [87]. Two-dimensional Majorana lattices with random sign couplings have additionally been numerically shown to support them [88]. Beyond this, thermal metal behaviour has also been seen in models of spinless p-wave topological superconductors by artificially inserting vortices in order to simulate temperatures above the Kosterlitz-Thouless transition point where these vortices proliferate [113]. However, in all of these examples the disorder is introduced by hand. In contrast, the case we have studied arises from the finite temperature physics of a single Hamiltonian, the Kitaev honeycomb spin model. The vortices that act as disorder to the fermionic degrees of freedom occur naturally as finite temperature excitations of this Hamiltonian.

In studying the thermal metal behaviour we focussed on temperatures well above the

vortex proliferation temperature. An interesting problem would be to study the onset of the thermal metal behaviour. This is hard, however, as the system sizes that are accessible to us are strongly limited by the numerical cost of the algorithm we use to generate thermal states. Simpler sampling schemes may be useful, such as generating random vortex patterns with a fixed average density corresponding to the thermal value. However, we should be careful using these approaches at lower vortex density since the Ising vortices are interacting [54]. These interactions will change the preferred spatial arrangements of the vortices at lower densities.

There is an interesting discrepancy between our results in section 6.4.3 and those of [105]. In that work the authors compute the thermal conductivity of the Kitaev honeycomb over a large range of temperatures and note that they observe no finite size dependence of the (longitudinal) thermal conductivity. Whereas, the thermal metal phase we have argued is present in the Ising regime at high temperatures is expected to have a thermal conductivity that diverges with system size. This divergence, however, is only logarithmic and so may difficult to observe. Alternatively, the method they employ to compute the conductivity might hide this behaviour. Their calculation involves using small size ($L = 12$) finite temperature gauge configurations as the unit cell of a larger system and computing the conductivity in the momentum space of that larger system. Thermal metals correspond to an infinite randomness renormalisation fixed point [106], which may make imposing a repeating length scale in this way problematic. It would be interesting to revisit the question of thermal conductivity and look closer for signs of the thermal metal behaviour.

7.

Conclusions

The focus of this thesis has been studying the effects of finite temperatures in different topological systems. First a topological quantum memory, followed by topological phases and finally a topologically ordered spin model. In this chapter I sum up by reviewing the results of chapters 3, 5 & 6, and suggesting future directions.

In chapter 3 we studied the coherence time τ of the toric code as a finite temperature quantum memory. The behaviour of τ was investigated by analysing the creation and diffusion of thermal excitations. Different processes occur on different timescales given by the rates obtained from the thermal master equation. These timescales were combined with arguments about the random walks of the excitations to give final estimates of τ. We compared these arguments to data from simulations and found good agreement.

Our primary interest in studying the toric code's coherence times was to compare the typical behaviour of the memory in a thermal environment to the previously demonstrated upper bounds [21, 23]. The scaling of these upper bounds follows the Arrhenius law applied the energy cost of quasiparticle creation (which is 2Δ since toric code anyons are created in pairs), explicitly $\tau \lesssim e^{2\beta\Delta}$. Studying the particle creation and diffusion processes allowed us to quantify the effects of entropy, which decrease the coherence time away from this upper bound. In the thermodynamic limit we found the entropic factors saturate giving a coherence time that scales as $\tau \sim e^{\beta\Delta}$. For small $(L \times L)$ system sizes, below a thermal length scale $L^* \sim e^{\beta\Delta/2}$, the coherence time lies between these limits $e^{\beta\Delta} \lesssim \tau \lesssim e^{2\beta\Delta}$. In this regime the entropic factors introduce a system size dependence to the coherence time $\tau \sim e^{2\beta\Delta} \ln({}^{L}/_{2})/L^2$, which decreases τ from the upper limit towards the lower limit with increasing system size.

An important innovation in the study of passive quantum memories was the discovery of partially self-correcting memories [17, 25, 26]. These models show it is possible to achieve better scaling of τ than the Arrhenius law applied to the energy cost of quasiparticle creation would imply. They still have a fixed cost to excite quasiparticles from the vacuum.

However, by engineering a complex, *glassy* energy landscape they manage to slow down the diffusion of quasiparticles. For example, the excitations of the *cubic code* have to transition through intermediate higher energy states in order to travel [25]. These diffusion times become the dominant component of the coherence time. For small system sizes, when decoherence is driven by the motion of one set of quasiparticles, coherence times then scale as $\tau \sim L^{\kappa\beta}$ [17, 25]. However, as the system becomes larger (relative to its own characteristic thermal length scale) entropy injects multiple sets of excitations. Decoherence then occurs when these different sets overlap and the coherence time no longer grows with L. In the thermodynamic limit the scaling of the coherence time saturates to $\tau \sim e^{\lambda\beta^2}$ [17, 25]. Understanding the interaction between energetic and entropic factors clearly will play an important role in identifying memories that replicate or improve on partial self-correction.

Chapter 5 studied thermal energy currents in Kitaev's honeycomb. For certain choices of parameters and vortex patterns, the fermionic degrees of freedom of the model carry topological superconductor phases with non-trivial Chern number [51, 53]. Bulk-boundary correspondence relates these bulk topological phases to conformal field theories (CFT) living on open boundaries of the system [90]. These CFTs should give rise to persistent and robust energy currents on the system's edge. Using the exact solution of the model we investigated these currents in a finite size system.

Energy current operators were defined from the energy continuity equation, giving local current observables that are expressed in terms of two-point fermionic correlations. We computed the edge current and studied its scaling with temperature, as well as how it is effected by the finite size ($L \times L$) of the system. In a particular temperature range, bounded from below by $1/L$ and from above by the bulk gap, the energy current was found to agree well with the predictions of CFT. We also confirmed that the currents are robust against small perturbations. That is not to say there were no deviations from the predictions of CFT however. We also noted some non-topological features of the currents, including a current at zero temperature whose value depends on the coupling parameters and system size. (CFT predicts a zero temperature current that scales with L, but in this case the currents we found did not quite match the expectations.) We attributed these

non-topological aspects to bulk contributions; all states of the model are chiral and the bulk states are not perfectly isolated from the edge.

One motivation to study these currents is the possibility of their direct observation in quantum simulations of Kitaev's honeycomb. Quantum simulator schemes with cold-atoms are already able to simulate other topological phases, in particular successfully replicating topological insulators [45]. There are proposals to extend those schemes to simulate the topological superconductor phases of Kitaev's honeycomb [114]. In such a setup it might be possible to directly measure the same current observables and compare the microscopic theory to the quantum simulation.

An interesting future direction for this work is to use the energy currents to probe the transitions between different topological phases of the Kitaev model. Currents and conductivities have yielded useful insights into the topological phases of integer quantum Hall systems. In particular numerical studies of the conductivity have been able to extract information about the quantum phase transitions between insulating phases and quantised conduction plateau [115, 116]. These can then be compared to abstract theory and experimental measurements. The details extracted include the critical exponents, allowing the universality class of the transition to be identified.

The goal of chapter 6 was to study the thermodynamics of the full Kitaev honeycomb spin model. We proceeded from the Kitaev fermionisation of the spins. This has the complication that the fermionic states of the model need to be projected into a particular subspace in order to be *physical*, meaning spin states. We first addressed this by outlining a procedure through which we can calculate finite temperature properties of the projected states of the fermions.

The physical projection was shown to reveal the topological degeneracy of the model in a simple way. This was carried out by inserting topologically non-trivial gauge loops into the ground state sector of the model, the no-vortex sector, and comparing the energy of the physical fermion states. We showed that the topological degeneracy of the ground state matches well to the abstract topological quantum field theory descriptions of the different phases. In the Ising anyon phase this is a three-fold degeneracy and in the toric code limit it is a four-fold degeneracy. We studied the finite size splitting that lifts these degeneracies

and found evidence that they vanish as the system becomes large. Additionally, we noted that the degeneracy depends on whether the number of lattice sites is even or odd. The degeneracies quoted above are for even numbers of lattice sites. For odd number of lattice sites the toric code has a two-fold degenerate ground state, and the Ising phase has a unique ground state. The even/odd behaviour we observe in the Ising phase has also been noted in the degeneracy of the Moore-Read state [100], which gives rise to Ising anyons in fractional quantum Hall systems.

The procedure to handle the fermions was combined with a Monte Carlo simulation that generates thermal distributions of vortices in order to study the thermodynamics of the full model. Several important features of the thermodynamics were already known from previous studies. In particular, the specific heat shows signatures of the fractionalisation of the spins [60]. This manifests as two distinct peaks in the specific heat, one associated with the vortices and the other with the Majorana fermions. We independently verified this behaviour in our data. By carefully accounting for the projection of the fermions, we were able to demonstrate that the thermal entropy at low temperatures reproduces the topological ground state degeneracies of the different regimes of the model. We investigated the behaviour of the fermion parity as a function of temperature and explained its behaviours by studying how the fermion gap varies. Finally, we studied the high temperature physics of the Ising regime and demonstrated that it behaves as a thermal metal.

The projection of the fermions also highlights some interesting thermal properties of the toric code limit of the model. In particular, since the fermions in the toric code regime do not become gapless, the projection continues to couple together the behaviour of the fermions and vortices even after the vortices have appeared to thermalise. Specifically, there is a range of temperatures (between the two specific heat peaks) where the vortices are disordered (the vortex density is half) but the total fermion parity remains even. The anyon fusion rules of the toric code imply these vortex configurations are not completely random, specifically there can only be even numbers of both e and m anyons. It would be interesting to probe deeper into the possible implications of this in a quantum error correction setting.

There are many interesting open directions of study. Thermal conductivities are particularly relevant as they reveal properties of the low energy topological phases and can

be probed experimentally. Other authors have begun investigating these properties in some parameter regimes [62, 63]. We would like to extend these studies and to probe the effect of fermion projection on the thermal response properties. In appendix C we begin to construct a framework for extending the description of the real-space current used in chapter 5 to compute dynamical responses through the *Kubo formula*.

Protecting quantum information from noise and errors is a difficult physics problem, as well as an engineering challenge, and has been an active area of research for over two decades [6]. In the coming years we can expect the small quantum devices being developed by tech giants such as IBM and Google to grow steadily larger, and as they do it may become possible to begin to incorporate some of the ideas developed in quantum error correction theory into these real world realisations. Conventionally quantum error correction has focussed on big picture question such as '*is it possible, even in principle, to protect quantum information from errors?*' (yes) or '*what is the most powerful code that can be designed, no matter how complicated?*' (currently, see [12] and [13]). While there are still many big open problems, increasingly feedback from experiments is beginning to reshape the topic around the constraints of real devices with questions such as '*given this real noise model and this limited number of qubits, what is the most effective error correction scheme I can implement?*'. Pivoting in this way will involve integrating the successes of quantum error correction with the tools of other fields, such as statistical mechanics, to better describe real situations and answer these hybrid questions. This is the spirit in which the work of chapter 3 was carried out.

Currently it remains an open problem to conclusively demonstrate the presence of anyons in real materials. Topological edge currents (chapter 5) and manifestations of topological properties in other thermodynamic responses (chapter 6) represent experimentally accessible stepping stones towards this goal. Eventual proof of the presence of anyonic excitations in a real material, besides being an incredible discovery of great fundamental importance, could eventually open the way to future generations of quantum devices fabricated from topologically ordered materials.

References

[1] M. H. Freedman, A. Kitaev, M. J. Larsen & Z. Wang, "Topological Quantum Computation," *Bulletin of the American Mathematical Society*, vol. 40, #1, 31–38, 2003.

[2] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy et al., "A blueprint for demonstrating quantum supremacy with superconducting qubits," *arXiv.org*, Sep. 2017, `1709.06678v1`.
Intel Press release, Oct. 2017, "Intel Delivers 17-Qubit Superconducting Chip with Advanced Packaging to QuTech", accessed on 30/01/2018.
W. Knight, "IBM Raises the Bar with a 50-Qubit Quantum Computer," Press coverage of IBM from MIT Technology Review, Nov. 2017, accessed on 30/01/2018.
W. Knight, "A Startup Uses Quantum Computing to Boost Machine Learning," Press coverage of Rigetti from MIT Technology Review, Dec. 2017, accessed on 31/01/2018.

[3] J. Wootton, "Why we need to make quantum games," Blog Post on medium, 2017, accessed on 30/01/2018.

[4] A. Y. Kitaev, "Fault-tolerant quantum computation by anyons," *Annals of Physics*, vol. 303, #1, 2–30, 2003.

[5] J. K. Pachos & S. H. Simon, "Focus on topological quantum computation," *New J. Phys.*, vol. 16, #6, 065003, 2014.

[6] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory," *Phys. Rev. A*, vol. 52, #4, R2493–R2496, Oct. 1995.

[7] P. W. Shor, "Fault-tolerant quantum computation," in "the Proceedings of the 37th Annual Symposium on Foundations of Computer Science," , 1996.

[8] A. M. Steane, "Error Correcting Codes in Quantum Theory," *Phys. Rev. Lett.*, vol. 77, #5, 793–797, Jul. 1996.

[9] J. Preskill, "Reliable quantum computers," *Proceedings of the Royal Society of London A*, vol. 454, #1969, 385–410, Jan. 1998.

[10] D. Gottesman, "Fault-Tolerant Quantum Computation with Constant Overhead," *arXiv:1310.02984*, vol. quant-ph, 2013.

[11] S. Bravyi, M. Suchara & A. Vargo, "Efficient algorithms for maximum likelihood decoding in the surface code," *Phys. Rev. A*, vol. 90, #3, 216, 2014.

[12] S. Bravyi & A. Cross, "Doubled Color Codes," *arXiv.org*, Sep. 2015, `1509.03239v1`.

[13] H. Bombin, "Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes," *New Journal of Physics*, vol. 17, #8, 083002, Aug. 2015.

[14] B. J. Brown, N. H. Nickerson & D. E. Browne, "Fault-tolerant error correction with the gauge color code," *Nat. Comms.*, vol. 7, 12302, Jul. 2016.

[15] F. Pastawski, B. Yoshida, D. Harlow & J. Preskill, "Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence," *Journal of High Energy Physics*, vol. 2015, #6, 163, Jun. 2015.

[16] B. M. Terhal, "Quantum error correction for quantum memories," *Rev. Mod. Phys.*, vol. 87, #2, 307–346, Apr. 2015.

[17] B. J. Brown, D. Loss, J. K. Pachos, C. N. Self & J. R. Wootton, "Quantum memories at finite temperature," *Rev. Mod. Phys.*, vol. 88, #4, 045005, 2016.

[18] D. Bacon, "Operator quantum error-correcting subsystems for self-correcting quantum memories," *Phys. Rev. A*, vol. 73, 012340, 2006.

[19] O. Landon-Cardinal & D. Poulin, "Local Topological Order Inhibits Thermal Stability in 2D," *Phys. Rev. Lett.*, vol. 110, #9, 090502, Feb. 2013.

[20] S. Bravyi & B. Terhal, "A no-go theorem for a two-dimensional self-correcting quantum memory based on stabilizer codes," *New J. Phys.*, vol. 11, #4, 043029, 2009.

[21] R. Alicki, M. Fannes & M. Horodecki, "On thermalization in Kitaev's 2D model," *J. Phys. A: Math. Theor.*, vol. 42, #6, 065303, Feb. 2009.

[22] K. Temme & M. J. Kastoryano, "How fast do stabilizer Hamiltonians thermalize?" *arXiv.org*, May 2015, `1505.07811v1`.

[23] A. Kómár, O. Landon-Cardinal & K. Temme, "Necessity of an energy barrier for self-correction of Abelian quantum doubles," *Phys. Rev. A*, vol. 93, #5, 28, 2016.

[24] K. Temme, "Thermalization Time Bounds for Pauli Stabilizer Hamiltonians," *Commun. Math. Phys.*, vol. 350, #2, 603–637, 2016.

[25] S. Bravyi & J. Haah, "Quantum Self-Correction in the 3D Cubic Code Model," *Phys. Rev. Lett.*, vol. 111, 200501, 2013.

[26] C. Castelnovo & C. Chamon, "Topological quantum glassiness," *Philosophical Magazine*, vol. 92, #1-3, 304–323, 2012.

[27] E. Dennis, A. Kitaev, A. Landahl & J. Preskill, "Topological quantum memory," *J. Math. Phys.*, vol. 43, 4452, 2002.

[28] M. B. Hastings, G. H. Watson & R. G. Melko, "Self-Correcting Quantum Memories Beyond the Percolation Threshold," *Phys. Rev. Lett.*, vol. 112, #7, 8, 2014.

[29] R. Alicki, M. Horodecki, P. Horodecki & R. Horodecki, "On Thermal Stability of Topological Qubit in Kitaev's 4D Model," *Open Systems & Information Dynamics*, vol. 17, #01, 1–20, 2010.

[30] C. G. Brell, "A proposal for self-correcting stabilizer quantum memories in 3 dimensions (or slightly less)," *New J. Phys.*, vol. 18, #1, 013050, 2016.

[31] K. P. Michnicki, "3D Topological Quantum Memory with a Power-Law Energy Barrier," *Phys. Rev. Lett.*, vol. 113, #13, 130501, 2014.

[32] B. J. Brown, A. Al-Shimary & J. K. Pachos, "Entropic Barriers for Two-Dimensional Quantum Memories," *Phys. Rev. Lett.*, vol. 112, #12, 120503, Mar. 2014.

[33] B. Yoshida, "Violation of arrhenius law below the transition temperature," *arXiv:1404.0457*, 2014.

[34] C. D. Freeman, C. M. Herdman, D. J. Gorman & K. B. Whaley, "Relaxation dynamics of the toric code in contact with a thermal reservoir: Finite-size scaling in a low-temperature regime," *Physical Review B*, vol. 90, #13, 134302, Oct. 2014.

[35] J. K. Pachos, *Introduction to Topological Quantum Computation*, Cambridge University Press, Cambridge, 2012.

[36] Y. Aharonov & D. Bohm, "Significance of electromagnetic potentials in the quantum theory," *Phys. Rev.*, vol. 115, #3, 485–491, 1959.

[37] C. Nayak, S. H. Simon, A. Stern, M. Freedman & S. Das Sarma, "Non-Abelian anyons and topological quantum computation," *Reviews of Modern Physics*, vol. 80, #3, 1083–1159, 2008.

[38] G. K. Brennen & J. K. Pachos, "Why should anyone care about computing with anyons?" *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 464, #2089, 1–24, 2008.

[39] X.-G. Wen, "Colloquium : Zoo of quantum-topological phases of matter," *Reviews of Modern Physics*, vol. 89, #4, 1064, 2017.

[40] V. Lahtinen & J. Pachos, "A Short Introduction to Topological Quantum Computation," *SciPost Physics*, vol. 3, #3, 2017.

[41] B. A. Bernevig & T. L. Hughes, *Topological Insulators and Topological Superconductors*, Princeton University Press, 2013.

[42] M. Z. Hasan & C. L. Kane, "Colloquium: Topological insulators," *Reviews of Modern Physics*, vol. 82, #4, 3045–3067, 2010.

[43] X.-L. Qi & S.-C. Zhang, "Topological insulators and superconductors," *Reviews of Modern Physics*, vol. 83, #4, 1057–1110, 2011.

[44] I. Bloch, J. Dalibard & S. Nascimbène, "Quantum simulations with ultracold quantum gases," *Nature Physics*, vol. 8, #4, 267–276, Apr. 2012.

[45] G. Jotzu, M. Messer, R. Desbuquois, M. Lebrat, T. Uehlinger, D. Greif et al., "Experimental realization of the topological Haldane model with ultracold fermions," *Nature*, vol. 515, #7526, 237–240, 2014.

[46] M. Aidelsburger, M. Lohse, C. Schweizer, M. Atala, J. T. Barreiro, S. Nascimbène et al., "Measuring the Chern number of Hofstadter bands with ultracold bosonic atoms," *Nature Physics*, vol. 11, #2, 162–166, 2014.

[47] V. J. Goldman & B. Su, "Resonant Tunneling in the Quantum Hall Regime: Measurement of Fractional Charge," *Science*, vol. 267, #5200, 1010–1012, Feb. 1995.

[48] R. de Picciotto, M. Reznikov, M. Heiblum, V. Umansky, G. Bunin & D. Mahalu, "Direct observation of a fractional charge," *Nature*, vol. 389, #6647, 162–164, 1997.

[49] R. L. Willett, L. N. Pfeiffer & K. W. West, "Measurement of filling factor 5/2 quasiparticle interference with observation of charge e/4 and e/2 period oscillations," *Proceedings of the National Academy of Sciences*, vol. 106, #22, 8853–8858, 2009.

[50] A. Banerjee, C. A. Bridges, J. Q. Yan, A. A. Aczel, L. Li, M. B. Stone et al., "Proximate Kitaev quantum spin liquid behaviour in a honeycomb magnet," *Nature Materials*, vol. 15, #7, 733–740, Jul. 2016.

[51] A. Kitaev, "Anyons in an exactly solved model and beyond," *Annals of Physics*, vol. 321, #1, 2–111, Jan. 2006.

[52] V. Lahtinen & J. K. Pachos, "Non-Abelian statistics as a Berry phase in exactly solvable models," *New Journal of Physics*, vol. 11, #9, 093027, 2009.

[53] V. Lahtinen & J. K. Pachos, "Topological phase transitions driven by gauge fields in an exactly solvable model," *Physical Review B*, vol. 81, #24, 245132, Jun. 2010.

[54] V. Lahtinen, A. W. W. Ludwig, J. K. Pachos & S. Trebst, "Topological liquid nucleation induced by vortex-vortex interactions in Kitaev's honeycomb model," *Physical Review B*, vol. 86, #7, 075115, Aug. 2012.

[55] C. N. Self, J. K. Pachos, J. R. Wootton & S. Iblisdir, "Conformal energy currents on the edge of a topological superconductor," *Physical Review B*, vol. 95, #11, 115141, Mar. 2017.

[56] G. Jackeli & G. Khaliullin, "Mott Insulators in the Strong Spin-Orbit Coupling Limit: From Heisenberg to a Quantum Compass and Kitaev Models," *Phys. Rev. Lett.*, vol. 102, #1, 017205, Jan. 2009.

[57] Y. Singh, S. Manni, J. Reuther, T. Berlijn, R. Thomale, W. Ku et al., "Relevance of the Heisenberg-Kitaev Model for the Honeycomb Lattice Iridates A_2IrO_3," *Phys. Rev. Lett.*, vol. 108, #12, 127203, Mar. 2012.

[58] Y. Kasahara, K. Sugii, T. Ohnishi, M. Shimozawa, M. Yamashita, N. Kurita et al., "Unusual thermal Hall effect in a Kitaev spin liquid candidate α-$RuCl_3$," *arXiv.org*, Sep. 2017, 1709.10286v2.

[59] J. Nasu & Y. Motome, "Thermodynamics of Chiral Spin Liquids with Abelian and Non-Abelian Anyons," *Phys. Rev. Lett.*, vol. 115, #8, 087203, Aug. 2015.

[60] J. Nasu, M. Udagawa & Y. Motome, "Thermal fractionalization of quantum spins in a Kitaev model: Temperature-linear specific heat and coherent transport of Majorana fermions," *Physical Review B*, vol. 92, #11, 115122, 2015.

[61] J. Nasu, J. Knolle, D. L. Kovrizhin, Y. Motome & R. Moessner, "Fermionic response from fractionalization in an insulating two-dimensional magnet," *Nature Physics*, vol. 12, #10, 912–915, 2016.

[62] J. Nasu, J. Yoshitake & Y. Motome, "Thermal Transport in the Kitaev Model," *Phys. Rev. Lett.*, vol. 119, #12, 127204, 2017.

[63] A. Metavitsiadis, A. Pidatella & W. Brenig, "Thermal transport in a two-dimensional Z2 spin liquid," *Physical Review B*, vol. 96, #20, 205121, 2017.

[64] S. Bravyi, M. Englbrecht, R. Koenig & N. Peard, "Correcting coherent errors with surface codes," *arXiv.org*, Oct. 2017, 1710.02270v1.

[65] S. Beale, J. Wallman, M. Gutiérrez, K. R. Brown & R. Laflamme, "Coherence in quantum error-correcting codes," *arXiv.org*, May 2018, 1805.08802v1.

[66] E. Huang, A. C. Doherty & S. Flammia, "Performance of quantum error correction with coherent errors," *arXiv.org*, May 2018, 1805.08227v1.

[67] S. Bravyi, M. B. Hastings & S. Michalakis, "Topological quantum order: Stability under local perturbations," *Journal of Mathematical Physics*, vol. 51, #9, 093512, 2010.

[68] H.-P. Breuer & F. Petruccione, "The Theory of Open Quantum Systems," Oxford University Press, 2007.

[69] E. B. Davies, "Markovian master equations," *Communications in Mathematical Physics*, vol. 39, #2, 91–110, 1974.

[70] A. Rivas, A. D. K. Plato, S. F. Huelga & M. B. Plenio, "Markovian master equations: a critical study," *New J. Phys.*, vol. 12, #11, 113032, 2010.

[71] A. J. Leggett, S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg & W. Zwerger, "Dynamics of the dissipative two-state system," *Reviews of Modern Physics*, vol. 59, #1, 1–85, Jan. 1987.

[72] W. Janke, Monte Carlo Methods in Classical Statistical Physics , vol. 739 of *Lecture Notes in Physics*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[73] A. B. Bortz, M. H. Kalos & J. L. Lebowitz, "A new algorithm for Monte Carlo simulation of Ising spin systems," *J. Comput. Phys.*, vol. 17, #1, 10–18, 1975.

[74] D. S. Wang, A. G. Fowler, C. D. Hill & L. C. L. Hollenberg, "Graphical algorithms and threshold error rates for the 2d colour code," *Quantum Information and Computation*, vol. 10, #9, 780–802, 2010.

[75] G. Duclos-Cianci & D. Poulin, "Fast Decoders for Topological Quantum Codes," *Phys. Rev. Lett.*, vol. 104, #5, 449, 2010.

[76] A. Hutter, J. R. Wootton & D. Loss, "Efficient Markov chain Monte Carlo algorithm for the surface code," *Phys. Rev. A*, vol. 89, #2, 0456, 2014.

[77] H. Anwar, B. J. Brown, E. T. Campbell & D. E. Browne, "Fast decoders for qudit topological codes," *New J. Phys.*, vol. 16, #6, 063038, 2014.

[78] J. W. Harrington, *Analysis of quantum error-correcting codes: symplectic lattice codes and toric codes*, Ph.D. thesis, California Institute of Technology, 2004.

[79] A. Hutter, D. Loss & J. R. Wootton, "Improved HDRG decoders for qudit and non-Abelian quantum error correction," *New Journal of Physics*, vol. 17, #3, 035017, Mar. 2015.

[80] J. K. Pachos & M. B. Plenio, "Three-Spin Interactions in Optical Lattices and Criticality in Cluster Hamiltonians," *Phys. Rev. Lett.*, vol. 93, #5, 056402, 2004.

[81] H. P. Büchler, A. Micheli & P. Zoller, "Three-body interactions with cold polar molecules," *Nature Physics*, vol. 3, #10, 726–731, 2007.

[82] H.-D. Chen & Z. Nussinov, "Exact results of the Kitaev model on a hexagonal lattice: spin states, string and brane correlators, and anyonic excitations," *J. Phys. A: Math. Theor.*, vol. 41, #7, 075001, Feb. 2008.

[83] H.-D. Chen & J. Hu, "Exact mapping between classical and topological orders in two-dimensional spin systems," *Physical Review B*, vol. 76, #19, 193101, Nov. 2007.

[84] X.-Y. Feng, G.-M. Zhang & T. Xiang, "Topological Characterization of Quantum Phase Transitions in a Spin- 1/2Model," *Phys. Rev. Lett.*, vol. 98, #8, 087204, Feb. 2007.

[85] F. L. Pedrocchi, S. Chesi & D. Loss, "Physical solutions of the Kitaev honeycomb model," *Physical Review B*, vol. 84, #16, 165414, 2011.

[86] E. H. Lieb, "Flux Phase of the Half-Filled Band," *Phys. Rev. Lett.*, vol. 73, #16, 2158–2161, Oct. 1994.

[87] V. Lahtinen, A. W. W. Ludwig & S. Trebst, "Perturbed vortex lattices and the stability of nucleated topological phases," *Physical Review B*, vol. 89, #8, 085121, Feb. 2014.

[88] C. R. Laumann, A. W. W. Ludwig, D. A. Huse & S. Trebst, "Disorder-induced Majorana metal in interacting non-Abelian anyon systems," *Physical Review B*, vol. 85, #16, 161301, Apr. 2012.

[89] Y. E. Kraus & A. Stern, "Majorana fermions on a disordered triangular lattice," *New Journal of Physics*, vol. 13, #10, 105006, Oct. 2011.

[90] G. Palumbo & J. K. Pachos, "Holographic correspondence in topological superconductors," *Annals of Physics*, vol. 372, 175–181, Sep. 2016.

[91] A. Cappelli, M. Huerta & G. R. Zemba, "Thermal transport in chiral conformal theories and hierarchical quantum Hall states," *Nuclear Physics B*, vol. 636, #3, 568–582, 2002.

[92] A. Furusaki, N. Nagaosa, K. Nomura & S. Ryu, "Electromagnetic and thermal responses in topological matter: topological terms, quantum anomalies and D-branes," *Comptes Rendus Physique*, vol. 14, #9-10, 871–883, 2013.

[93] H. Sumiyoshi & S. Fujimoto, "Quantum Thermal Hall Effect in a Time-Reversal-Symmetry-Broken Topological Superconductor in Two Dimensions: Approach from Bulk Calculations," *J. Phys. Soc. Jpn.*, vol. 82, #2, 023602, 2013.

[94] P. Di Francesco, P. Mathieu & D. Senechal, *Conformal Field Theory*, Graduate Texts in Contemporary Physics, Springer-Verlag, New York, 1997.

[95] K. Meichanetzidis, M. Cirio, J. K. Pachos & V. Lahtinen, "Anatomy of fermionic entanglement and criticality in Kitaev spin liquids," *Physical Review B*, vol. 94, #11, 115158, Sep. 2016.

[96] G. Moore & N. Read, "Nonabelions in the fractional quantum hall effect," *Nuclear Physics B*, vol. 360, #2-3, 362–396, Aug. 1991.

[97] B. I. Halperin, "Quantized Hall conductance, current-carrying edge states, and the existence of extended states in a two-dimensional disordered potential," *Physical Review B*, vol. 25, #4, 2185–2190, Feb. 1982.

[98] J. Nasu, M. Udagawa & Y. Motome, "Thermal fractionalization of quantum spins in a Kitaev model: Temperature-linear specific heat and coherent transport of Majorana fermions," *Physical Review B*, vol. 92, #11, 115122, Sep. 2015.

[99] J. Nasu & Y. Motome, "Thermodynamics of Chiral Spin Liquids with Abelian and Non-Abelian Anyons," *Phys. Rev. Lett.*, vol. 115, #8, 087203, Aug. 2015.

[100] N. Read & D. Green, "Paired states of fermions in two dimensions with breaking of parity and time-reversal symmetries and the fractional quantum Hall effect," *Physical Review B*, vol. 61, #15, 10267–10297, Apr. 2000.

[101] G. Kells, A. T. Bolukbasi, V. Lahtinen, J. K. Slingerland, J. K. Pachos & J. Vala, "Topological degeneracy and vortex manipulation in Kitaev's honeycomb model," *Phys. Rev. Lett.*, vol. 101, 240404, 2008.

[102] G. Kells, J. K. Slingerland & J. Vala, "Description of Kitaev's honeycomb model with toric-code stabilizers," *Physical Review B*, vol. 80, #12, 125415, 2009.

[103] G. Kells, D. Sen, J. K. Slingerland & S. Vishveshwara, "Topological blocking in quantum quench dynamics," *Physical Review B*, vol. 89, #23, 369, 2014.

[104] D. J. Earl & M. W. Deem, "Parallel tempering: Theory, applications, and new perspectives," *Phys. Chem. Chem. Phys.*, vol. 7, #23, 3910–3916, Nov. 2005.

[105] J. Nasu, J. Yoshitake & Y. Motome, "Thermal Transport in the Kitaev Model," *Phys. Rev. Lett.*, vol. 119, 127204, Sep. 2017.

[106] M. Bocquet, D. Serban & M. R. Zirnbauer, "Disordered 2d quasiparticles in class D: Dirac fermions with random mass, and dirty superconductors," *Nuclear Physics B*, vol. 578, #3, 628–680, 2000.

[107] F. Evers & A. D. Mirlin, "Anderson transitions," *Rev. Mod. Phys.*, vol. 80, 1355–1417, Oct. 2008.

[108] A. Altland & M. R. Zirnbauer, "Nonstandard symmetry classes in mesoscopic normal-superconducting hybrid structures," *Physical Review B*, vol. 55, 1142–1161, Jan. 1997.

[109] A. Mildenberger, F. Evers, A. D. Mirlin & J. T. Chalker, "Density of quasiparticle states for a two-dimensional disordered system: Metallic, insulating, and critical behavior in the class-D thermal quantum Hall effect," *Physical Review B*, vol. 75, 245321, Jun. 2007.

[110] A. Metavitsiadis, A. Pidatella & W. Brenig, "Thermal transport in a two-dimensional Z2spin liquid," *Physical Review B*, vol. 96, #20, 205121, Nov. 2017.

[111] E. P. Wigner, "On a class of analytic functions from the quantum theory of collisions," in "The Collected Works of Eugene Paul Wigner," Springer, 1993.

[112] V. I. Fal'ko & K. B. Efetov, "Multifractality: Generic Property of Eigenstates of 2D Disordered Metals," *Europhysics Letters (EPL)*, vol. 32, #8, 627–632, Jul. 2007.

[113] B. Bauer, R. M. Lutchyn, M. B. Hastings & M. Troyer, "Effect of thermal fluctuations in topological p-wave superconductors," *Physical Review B*, vol. 87, 014503, Jan. 2013.

[114] A. Micheli, G. K. Brennen & P. Zoller, "A toolbox for lattice-spin models with polar molecules," *Nature Physics*, vol. 2, #5, 341–347, 2006.

[115] J. Song & E. Prodan, "Characterization of the quantized Hall insulator phase in the quantum critical regime," *EPL (Europhysics Letters)*, vol. 105, #3, 37001, 2014.

[116] J. Priest, S. P. Lim & D. N. Sheng, "Scaling behavior of the insulator-to-plateau transition in a topological band model," *Physical Review B*, vol. 89, #16, 165422, Apr. 2014.

[117] E. Jones, T. Oliphant, P. Peterson et al., "SciPy: Open source scientific tools for Python," , 2001–, `http://www.scipy.org/`.

[118] F. Perez & B. E. Granger, "IPython: A System for Interactive Scientific Computing," *Computing in Science & Engineering*, vol. 9, #3, 21–29, 2007.

[119] J. D. Hunter, "Matplotlib: A 2D Graphics Environment," *Comput. Sci. Eng.*, vol. 9, #3, 90–95, 2007.

[120] L. Dalcín, R. Paz, M. Storti & J. D'Elía, "MPI for Python: Performance improvements and MPI-2 extensions," *Journal of Parallel and Distributed Computing*, vol. 68, #5, 655–662, 2008.

[121] T. Oliphant, "Speeding up Python (NumPy, Cython, and Weave)," Blog Post, 2011.

[122] C. J. Geyer, "Practical Markov Chain Monte Carlo," *Statistical Science*, vol. 7, #4, 473–483, Nov. 1992.

[123] M. Weigel & W. Janke, "Error estimation and reduction with cross correlations," *Phys. Rev. E*, vol. 81, #6, 066701, Jun. 2010.

[124] B. Efron, "Computers and the Theory of Statistics: Thinking the Unthinkable," *SIAM Review*, vol. 21, #4, 460–480, Oct. 1979.

[125] M. E. G. Newman & G. T. Barkema, *Monte Carlo methods in statistical physics*, Clarendon Press, Oxford, 1999.

[126] H. G. Katzgraber, "Introduction to Monte Carlo Methods," *arXiv.org*, May 2009, `0905.1629v3`.

[127] J. D. Chodera, "A Simple Method for Automated Equilibration Detection in Molecular Simulations," *J. Chem. Theory Comput.*, vol. 12, #4, 1799–1805, 2016.

[128] J. Machta, "Strengths and weaknesses of parallel tempering," *Phys. Rev. E*, vol. 80, #5, 056706, Nov. 2009.

[129] H. G. Katzgraber, S. Trebst, D. A. Huse & M. Troyer, "Feedback-optimized parallel tempering Monte Carlo," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2006, #03, P03018–P03018, 2006.

[130] J. M. Luttinger, "Theory of Thermal Transport Coefficients," *Phys. Rev.*, vol. 135, #6A, A1505–A1514, 1964.

[131] D. J. Scalapino, S. R. White & S. Zhang, "Insulator, metal, or superconductor: The criteria," *Physical Review B*, vol. 47, #13, 7995–8007, Apr. 1993.

[132] R. Steinigeweg, J. Gemmer & W. Brenig, "Spin and energy currents in integrable and nonintegrable spin-1/2 chains: A typicality approach to real-time autocorrelations," *Physical Review B*, vol. 91, #10, 104404, Mar. 2015.

[133] P. D. Rowell, "Signal Processing: Continuous and Discrete," *Massachusetts Institute of Technology: MIT OpenCourseWare*, `https://ocw.mit.edu/`, Lecture Notes: MIT Course Number 2.161, accessed on 30/03/2017.

Appendices

A.

Kitaev's Honeycomb
Solved in a Matrix Language

Here the constructions in chapter 4 are repeated using an entirely matrix based presentation. This way of approaching the model helped me to understand it better, coming from a quantum information background.

This appendix mixes operators with matrices and vectors, so to be concretely clear about the nature of each object I have used the following conventions:

- $\hat{c}$, a hat indicates an operator
- $\mathbf{A}$, bold font denotes a matrix
  - $\hat{\mathbf{M}}$, combining the bold font with the hat means this is a matrix whose elements are operators
- $\underline{v}$, underlining means a vector
  - by convention $\underline{v}$ will always be a column vector, and $\underline{v}^T$ (or $\underline{v}^\dagger$ for complex vectors) the row vector
  - $\underline{\hat{w}}$, again combining the underline with the hat means the elements of this vector are operators

With the vortex sector specified and the gauge configuration fixed we can write the Hamiltonian of the fermionic degrees of freedom, eq. (4.8) in the main text, as

$$\begin{aligned}\hat{H} &= \frac{i}{4}\sum_{i,j=1}^{N} A_{ij}\hat{c}_i\hat{c}_j \\ &= \frac{1}{4}\left(\hat{c}_1 \dots \hat{c}_N\right)\,(i\mathbf{A})\begin{pmatrix}\hat{c}_1\\ \vdots \\ \hat{c}_N\end{pmatrix} \\ &= \frac{1}{4}\,\underline{\hat{c}}^T\,(i\mathbf{A})\,\underline{\hat{c}} \quad . \end{aligned} \tag{A.1}$$

Written in this form it is clear the task of diagonalising $\hat{H}$ is equivalent to diagonalising the hermitian matrix of numbers $i\mathbf{A}$. To do this we employ some standard results for real anti-symmetric (also called skew-symmetric) matrices.

A.1. Skew-symmetric matrices

For any skew-symmetric matrix $\mathbf{A}$ the eigenvalues are purely imaginary and occur in plus/minus pairs $\pm i\varepsilon_k$, for $k = 1, 2, \ldots, {}^{N}/_{2}$. The real valued matrix can therefore be diagonalised by complex matrices

$$\mathbf{A} = \mathbf{U}\,\mathbf{D}\,\mathbf{U}^{\dagger} = (\,\underline{u}_1^{+},\,\underline{u}_1^{-},\,\underline{u}_2^{+},\,\underline{u}_2^{-},\,\ldots\,)\begin{pmatrix} i\varepsilon_1 & 0 & 0 & 0 & \ldots \\ 0 & -i\varepsilon_1 & 0 & 0 & \ldots \\ 0 & 0 & i\varepsilon_2 & 0 & \ldots \\ 0 & 0 & 0 & -i\varepsilon_2 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}\begin{pmatrix} (\underline{u}_1^{+})^{\dagger} \\ (\underline{u}_1^{-})^{\dagger} \\ (\underline{u}_2^{+})^{\dagger} \\ (\underline{u}_2^{-})^{\dagger} \\ \vdots \end{pmatrix}$$

where $\mathbf{U}$ is a unitary matrix. Limiting ourselves to purely real matrices, it is not possible to fully diagonalise $\mathbf{A}$ but it can be put in block diagonal form

$$\mathbf{A} = \mathbf{Q}\,\mathbf{\Sigma}\,\mathbf{Q}^{T} = (\,\underline{q}_1,\,\underline{q}_2,\,\underline{q}_3,\,\underline{q}_4,\,\ldots\,)\begin{pmatrix} 0 & \varepsilon_1 & 0 & 0 & \ldots \\ -\varepsilon_1 & 0 & 0 & 0 & \ldots \\ 0 & 0 & 0 & \varepsilon_2 & \ldots \\ 0 & 0 & -\varepsilon_2 & 0 & \ldots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}\begin{pmatrix} \underline{q}_1^{T} \\ \underline{q}_2^{T} \\ \underline{q}_3^{T} \\ \underline{q}_4^{T} \\ \vdots \end{pmatrix}$$

where $\mathbf{Q}$ is a real orthogonal matrix. By relating these two forms of $\mathbf{A}$ we can demonstrate an important property of the eigenvectors. Namely that the complex eigenvectors $\underline{u}_k^{+}$ and $\underline{u}_k^{-}$, corresponding to eigenvalues pair $\pm i\varepsilon_k$, are related to each other by complex conjugation $\underline{u}_k^{-} = (\underline{u}_k^{+})^{*}$.

Consider the simpler case of σ_z and σ_y. We can relate the two using σ_y's eigenvectors

packaged in the form of unitary matrix $\mathbf{W}$

$$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ i & -i \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix} = \mathbf{W}\, \sigma_z\, \mathbf{W}^\dagger \,.$$

Up to a factor of i this is identical to what is happening in the diagonal blocks of $\mathbf{D}$ and $\mathbf{\Sigma}$ so, using the ${}^{N}\!/_{2} \times {}^{N}\!/_{2}$ identity matrix $\mathbb{1}_{N/2}$ to distribute $\mathbf{W}$ onto each block, we can write the same relation

$$\mathbf{\Sigma} = (\mathbb{1}_{N/2} \otimes \mathbf{W})\, \mathbf{D}\, (\mathbb{1}_{N/2} \otimes \mathbf{W}^\dagger) \,.$$

We can go further and use this to relate the columns of $\mathbf{U}$ and $\mathbf{Q}$,

$$\begin{aligned} (\,\underline{u}_1^+,\, \underline{u}_1^-,\, \underline{u}_2^+,\, \underline{u}_2^-,\, \dots\,) &= (\,\underline{q}_1,\, \underline{q}_2,\, \underline{q}_3,\, \underline{q}_4,\, \dots\,)\, (\mathbb{1}_{N/2} \otimes \mathbf{W}) \\ &= (\,\underline{q}_1,\, \underline{q}_2,\, \underline{q}_3,\, \underline{q}_4,\, \dots\,)\, \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 & 0 & 0 & \dots \\ i & -i & 0 & 0 & \dots \\ 0 & 0 & 1 & 1 & \dots \\ 0 & 0 & i & -i & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \\ &= \tfrac{1}{\sqrt{2}} \left(\underline{q}_1 + i\underline{q}_2,\ \underline{q}_1 - i\underline{q}_2,\ \underline{q}_3 + i\underline{q}_4,\ \underline{q}_3 - i\underline{q}_4,\, \dots \right) . \end{aligned} \tag{A.2}$$

Reading off from this we see that $\underline{u}_k^- = (\underline{u}_k^+)^*$ (as with the eigenvectors of σ_y). From now we will drop the $\pm$ superscript and write $\underline{u}_k^+ = \underline{u}_k$ and $\underline{u}_k^- = \underline{u}_k^*$. Now let us return to diagonalising $\hat{H}$.

A.2. Normal modes

Using the results of the previous section we diagonalise $\hat{H}$ to find the normal modes of the system. From equation (A.1) and the different forms of $\mathbf{A}$ we have

$$\hat{H} \;=\; \frac{1}{4}\, \underline{\hat{c}}^T\, (i\mathbf{A})\, \underline{\hat{c}} \;=\; \frac{1}{4}\, \underline{\hat{c}}^T\, \mathbf{Q}\, (i\mathbf{\Sigma})\, \mathbf{Q}^T\, \underline{\hat{c}} \;=\; \frac{1}{4}\, \underline{\hat{c}}^T\, \mathbf{U}\, (i\mathbf{D})\, \mathbf{U}^\dagger\, \underline{\hat{c}} \,. \tag{A.3}$$

These two different forms of $\mathbf{A}$, in terms of $\mathbf{Q}$ & $\mathbf{\Sigma}$ or $\mathbf{U}$ & $\mathbf{D}$, give the Majorana and associated complex fermion representations of the normal modes.

First consider multiplying out the new operators defined by $\mathbf{U}$ to obtain their expression in terms of the original Majorana,

$$\begin{pmatrix} \hat{f}_1 \\ \hat{g}_1 \\ \vdots \end{pmatrix} = \frac{1}{\sqrt{2}}\, \mathbf{U}^\dagger \begin{pmatrix} \hat{c}_1 \\ \hat{c}_2 \\ \vdots \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \underline{u}_1^\dagger \\ (\underline{u}_1^*)^\dagger \\ \vdots \end{pmatrix} \begin{pmatrix} \hat{c}_1 \\ \hat{c}_2 \\ \vdots \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \underline{u}_1 \cdot \underline{\hat{c}} \\ \underline{u}_1^* \cdot \underline{\hat{c}} \\ \vdots \end{pmatrix} \tag{A.4}$$

where $a \cdot b$ denotes the vector dot product. The extra factor of $1/\sqrt{2}$ is added for normalisation (so that the eigenvalues of $\hat{f}_k^\dagger \hat{f}_k$ are between zero and one). We see that as well as being complex, due of the elements of $\mathbf{U}$, the fermionic operators $\hat{f}_1$ and $\hat{g}_1$ are related by $\hat{g}_1 = \hat{f}_1^\dagger$. Multiplying out equation (A.3) in terms of these new fermions

$$\begin{aligned} \hat{H} &= \frac{1}{2} \left(\hat{f}_1^\dagger,\, \hat{f}_1,\, \hat{f}_2^\dagger,\, \hat{f}_2,\, \dots \right) \begin{pmatrix} -\varepsilon_1 & 0 & 0 & 0 & \dots \\ 0 & \varepsilon_1 & 0 & 0 & \dots \\ 0 & 0 & -\varepsilon_2 & 0 & \dots \\ 0 & 0 & 0 & \varepsilon_2 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} \hat{f}_1 \\ \hat{f}_1^\dagger \\ \hat{f}_2 \\ \hat{f}_2^\dagger \\ \vdots \end{pmatrix} \\ &= \frac{1}{2} \sum_{k=1}^{N/2} \left(-\varepsilon_k\, \hat{f}_k^\dagger \hat{f}_k + \varepsilon_k\, \hat{f}_k \hat{f}_k^\dagger \right) . \end{aligned}$$

The operator $\hat{f}_k^\dagger \hat{f}_k$ detects the presence of a particle in mode k by projecting onto the filled state and $\hat{f}_k \hat{f}_k^\dagger$ projects onto the empty state, detecting a hole. So we can read this as the occupation of normal mode k contributes energy $-\varepsilon_k/2$, and mode k being empty contributes energy $+\varepsilon_k/2$. This sounds a little backwards and illustrates part of the particle-hole symmetry of the spectrum. We arbitrarily picked the operators $\hat{f}$ rather than $\hat{g}$ and this gave us a *hole-basis* where exciting the normal modes removes energy from the system, doing the opposite would have instead given us a *particle-basis* where quasi-particle excitations carry positive energy,

$$\hat{H} = \sum_{k=1}^{N/2} (-\varepsilon_k) \left(\hat{f}_k^\dagger \hat{f}_k - \tfrac{1}{2} \right) = \sum_{k=1}^{N/2} \varepsilon_k \left(\hat{g}_k^\dagger \hat{g}_k - \tfrac{1}{2} \right) . \tag{A.5}$$

Similarly we can multiply out the operators defined by $\mathbf{Q}$ to get another representation of the normal modes

$$\begin{pmatrix} \hat{b}'_1 \\ \hat{b}''_1 \\ \vdots \end{pmatrix} = \mathbf{Q}^T \begin{pmatrix} \hat{c}_1 \\ \hat{c}_2 \\ \vdots \end{pmatrix} = \begin{pmatrix} \underline{q}_1^T \\ \underline{q}_2^T \\ \vdots \end{pmatrix} \begin{pmatrix} \hat{c}_1 \\ \hat{c}_2 \\ \vdots \end{pmatrix} = \begin{pmatrix} \underline{q}_1 \cdot \underline{\hat{c}} \\ \underline{q}_2 \cdot \underline{\hat{c}} \\ \vdots \end{pmatrix} . \tag{A.6}$$

Since $\mathbf{Q}$ is a real matrix these new operators are also real, and so these fermions are still Majorana fermions. Multiplying out the Hamiltonian in this representation gives

$$\begin{aligned} \hat{H} &= \frac{1}{4} \left(\hat{b}'_1, \hat{b}''_1, \hat{b}'_2, \hat{b}''_2, \dots \right) \begin{pmatrix} 0 & i\varepsilon_1 & 0 & 0 & \dots \\ -i\varepsilon_1 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & i\varepsilon_2 & \dots \\ 0 & 0 & -i\varepsilon_2 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} \hat{b}'_1 \\ \hat{b}''_1 \\ \hat{b}'_2 \\ \hat{b}''_2 \\ \vdots \end{pmatrix} \\ &= \frac{1}{4} \sum_{k=1}^{N/2} \left(-i\varepsilon_k \, \hat{b}'_k \hat{b}''_k + i\varepsilon_k \, \hat{b}''_k \hat{b}'_k \right) = \frac{i}{2} \sum_{k=1}^{N/2} \varepsilon_k \, \hat{b}''_k \hat{b}'_k \, . \end{aligned}$$

Combining equations (A.4) and (A.6) with (A.2) we can relate the diagonal fermion creation and annihilation operators $\hat{g}$ and $\hat{g}^\dagger$ to the Majorana normal modes $\hat{b}'$ and $\hat{b}''$ getting the expected expressions,

$$\hat{g}_k = \tfrac{1}{2} \left(\hat{b}'_k - i\hat{b}''_k \right) , \quad \hat{g}_k^\dagger = \tfrac{1}{2} \left(\hat{b}'_k + i\hat{b}''_k \right) .$$

A.3. Correlation functions

An important set of quantities to characterise the properties of fermionic states, with density operator $\hat{\rho}$, are the two-point correlation functions

$$\langle \hat{c}_i \hat{c}_j \rangle = \mathrm{tr}\left(\hat{\rho}\, \hat{c}_i \hat{c}_j\right) .$$

We can evaluate these expectation values by inverting the transformation to energy eigenmodes, equation (A.4). This gives an expression

$$\begin{pmatrix} \hat{c}_1 \\ \hat{c}_2 \\ \vdots \end{pmatrix} = \sqrt{2}\, \mathbf{U} \begin{pmatrix} \hat{g}_1^\dagger \\ \hat{g}_1 \\ \vdots \end{pmatrix} . \tag{A.7}$$

Taking the outer product of $\underline{\hat{c}}$ with $\underline{\hat{c}}^T$ gives a matrix of correlation operators, as well as the transformation of these operators from a spatial basis to the diagonal basis

$$\begin{pmatrix} \hat{c}_1\hat{c}_1 & \hat{c}_1\hat{c}_2 & \hat{c}_1\hat{c}_3 & \hat{c}_1\hat{c}_4 & \dots \\ \hat{c}_2\hat{c}_1 & \hat{c}_2\hat{c}_2 & \hat{c}_2\hat{c}_3 & \hat{c}_2\hat{c}_4 & \dots \\ \hat{c}_3\hat{c}_1 & \hat{c}_3\hat{c}_2 & \hat{c}_3\hat{c}_3 & \hat{c}_3\hat{c}_4 & \dots \\ \hat{c}_4\hat{c}_1 & \hat{c}_4\hat{c}_2 & \hat{c}_4\hat{c}_3 & \hat{c}_4\hat{c}_4 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} = 2\,\mathbf{U} \begin{pmatrix} \hat{g}_1^\dagger \\ \hat{g}_1 \\ \hat{g}_2^\dagger \\ \hat{g}_2 \\ \vdots \end{pmatrix} \left(\hat{g}_1,\, \hat{g}_1^\dagger,\, \hat{g}_2,\, \hat{g}_2^\dagger \dots\right) \mathbf{U}^\dagger$$

$$\begin{pmatrix} \mathbb{1} & \hat{c}_1\hat{c}_2 & \hat{c}_1\hat{c}_3 & \hat{c}_1\hat{c}_4 & \dots \\ \hat{c}_2\hat{c}_1 & \mathbb{1} & \hat{c}_2\hat{c}_3 & \hat{c}_2\hat{c}_4 & \dots \\ \hat{c}_3\hat{c}_1 & \hat{c}_3\hat{c}_2 & \mathbb{1} & \hat{c}_3\hat{c}_4 & \dots \\ \hat{c}_4\hat{c}_1 & \hat{c}_4\hat{c}_2 & \hat{c}_4\hat{c}_3 & \mathbb{1} & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} = 2\,\mathbf{U} \begin{pmatrix} \hat{g}_1^\dagger\hat{g}_1 & 0 & \hat{g}_1^\dagger\hat{g}_2 & \hat{g}_1^\dagger\hat{g}_2^\dagger & \dots \\ 0 & \hat{g}_1\hat{g}_1^\dagger & \hat{g}_1\hat{g}_2 & \hat{g}_1\hat{g}_2^\dagger & \dots \\ \hat{g}_2^\dagger\hat{g}_1 & \hat{g}_2^\dagger\hat{g}_1^\dagger & \hat{g}_2^\dagger\hat{g}_2 & 0 & \dots \\ \hat{g}_2\hat{g}_1 & \hat{g}_2\hat{g}_1^\dagger & 0 & \hat{g}_2\hat{g}_2^\dagger & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \mathbf{U}^\dagger .$$

We want to evaluate the expectation values of the elements of this matrix, $\langle \hat{c}_i \hat{c}_j \rangle$. If we carried out the multiplication on the right hand side of the expression each element of the resultant matrix would be a linear sum of the $\underline{\hat{g}}\,\underline{\hat{g}}^\dagger$ elements weighted by the complex numbers in the matrix $\mathbf{U}$. We can then see that to evaluate the $\langle \hat{c}_i \hat{c}_j \rangle$ we may replace

each element of $\underline{\hat{g}}\,\underline{\hat{g}}^{\dagger}$ by its expectation value, i.e.

$$\begin{pmatrix} 1 & \langle \hat{c}_1 \hat{c}_2 \rangle & \langle \hat{c}_1 \hat{c}_3 \rangle & \langle \hat{c}_1 \hat{c}_4 \rangle & \dots \\ \langle \hat{c}_2 \hat{c}_1 \rangle & 1 & \langle \hat{c}_2 \hat{c}_3 \rangle & \langle \hat{c}_2 \hat{c}_4 \rangle & \dots \\ \langle \hat{c}_3 \hat{c}_1 \rangle & \langle \hat{c}_3 \hat{c}_2 \rangle & 1 & \langle \hat{c}_3 \hat{c}_4 \rangle & \dots \\ \langle \hat{c}_4 \hat{c}_1 \rangle & \langle \hat{c}_4 \hat{c}_2 \rangle & \langle \hat{c}_4 \hat{c}_3 \rangle & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} = 2\,\mathbf{U} \begin{pmatrix} \langle \hat{g}_1^{\dagger} \hat{g}_1 \rangle & 0 & \langle \hat{g}_1^{\dagger} \hat{g}_2 \rangle & \langle \hat{g}_1^{\dagger} \hat{g}_2^{\dagger} \rangle & \dots \\ 0 & \langle \hat{g}_1 \hat{g}_1^{\dagger} \rangle & \langle \hat{g}_1 \hat{g}_2 \rangle & \langle \hat{g}_1 \hat{g}_2^{\dagger} \rangle & \dots \\ \langle \hat{g}_2^{\dagger} \hat{g}_1 \rangle & \langle \hat{g}_2^{\dagger} \hat{g}_1^{\dagger} \rangle & \langle \hat{g}_2^{\dagger} \hat{g}_2 \rangle & 0 & \dots \\ \langle \hat{g}_2 \hat{g}_1 \rangle & \langle \hat{g}_2 \hat{g}_1^{\dagger} \rangle & 0 & \langle \hat{g}_2 \hat{g}_2^{\dagger} \rangle & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \mathbf{U}^{\dagger}\,.$$

It will usually be the case that all off-diagonal elements of the matrix of $\hat{g}$ operators vanish as they represent coherences between different diagonal modes.

B.

Numerical notes and methods

B.1. Kitaev's Honeycomb

The following are notes on some practical elements of the numerical work I carried out on the Kitaev Honeycomb. The numerics were written in python using the NumPy and SciPy libraries [117]. Notebook work, such as data analysis, was carried out in IPython [118] and Matplotlib was used to generate figures [119]. Additionally, mpi4py [120] is used to implement parallel tempering via distributed parallel computing. Unless otherwise referenced or defined, the notation used in this section matches appendix A.

B.1.1. Obtaining diagonal modes

Diagonalising $i\mathbf{A}$ with complex numbers using SciPy's `linalg.eigh` seems to be problematic. Specifically, the returned matrix of eigenvectors $\mathbf{U}$ do not satisfy the physical requirement that the two-point correlations $\langle ic_j c_k \rangle$ should be purely real. This requirement is separate to just finding $\mathbf{U}$ and $\mathbf{D}$ such that $i\mathbf{A} = \mathbf{U}\,(i\mathbf{D})\,\mathbf{U}^\dagger$ and the function failing to satisfy this additional condition arises from the fact that we are using complex numbers to try to solve a problem that only really needs real numbers.

This physical condition is satisfied by instead directly computing the real-valued block-diagonal form $\mathbf{A} = \mathbf{Q}\,\mathbf{\Sigma}\,\mathbf{Q}^T$, calling Scipy's `linalg.schur` function with the call:

```
Σ, Q  = scipy.linalg.schur(A, output='real', lwork=None, overwrite_a=True,
sort=None, check_finite=True)
```

A disadvantage of this method is that the blocks are not sorted by energies. (The option `sort=None` is set in the function call since degeneracies can otherwise cause errors.) The eigenvalues and eigenvectors can be sorted separately if needed. Importantly, the $\mathbf{\Sigma}$ and $\mathbf{Q}$ obtained are still not in the form we expect them, as discussed in the next note.

B.1.2. Getting Q in *canonical* form

The physical fermion parity is determined by $\prod_{(i,j)} u_{ij}\,(\det \mathbf{Q})$, eq. (6.2). But to reliably and meaningfully compute $\det \mathbf{Q}$ we need $\mathbf{Q}$ and $\mathbf{\Sigma}$ to be in the form given on page 124, which is not guaranteed by `scipy.linalg.schur` (`sls`).

A particular problem is that `sls` does not arrange the positive and negative eigenvalues of the blocks of $\mathbf{\Sigma}$ consistently. Sometimes the positive eigenvalue is in the upper right (as we want) but sometimes it is in the lower left, and additionally it changes from block to block. If the eigenvalues are in *non-canonical* arrangement then the related columns of $\mathbf{Q}$ need to be swapped, which multiplies $\det \mathbf{Q}$ by -1. Another problem is if the spectrum contains zero modes these are not paired into 2×2 blocks in $\mathbf{\Sigma}$ by `sls`. Instead they appear separately as single, randomly positioned, entries on the diagonal.

Both of these issues are solved with the following code snippet, which puts $\mathbf{Q}$ and $\mathbf{\Sigma}$ into the canonical form of page 124. However, this fix is not robust against exact degeneracies in the spectrum. We must first lift the degeneracy slightly by adding a small amount of random local noise to the Hamiltonian.

```
# add small random noise to lift degeneracies
R = 2.*(numpy.random.rand(2*Lx*Ly,2*Lx*Ly) - 0.5)
R = R + np.transpose(R)
dA = 1E-6*numpy.sign(A)*R
# get the block-diagonal fully-real Schur normal form
Sigma,Q = scipy.linalg.schur(A + dA, output='real', lwork=None,\
    overwrite_a=True, sort=None, check_finite=True)

# flatten columns of Sigma to get list of eigenvalues
eig_vals = numpy.sum(Sigma,axis=0)
# get the -ve to +ve ordering of the eigenvalues
permutation = numpy.argsort(eig_vals)
# fold permutation back onto itself to get correct
# ordering for columns of Q matrix
q_permutation = numpy.zeros(2*Lx*Ly,dtype='i')
q_permutation[::2] = permutation[:2*Lx*Ly//2]
q_permutation[1::2] = permutation[-1:-2*Lx*Ly//2-1:-1]

# rearrange Q and Sigma into canonical form
Q = Q[:,q_permutation]
Sigma = Sigma[:,q_permutation]
Sigma = Sigma[q_permutation,:]
```

A test to check this is working is that $\prod_{(i,j)} u_{ij}\,(\det \mathbf{Q})$ should be invariant under local gauge transformations [85]. That is, it should not change when we insert any gauge loop

with trivial topology.

B.1.3. Robustly computing projected partition functions

Iteratively computing the partition function directly is numerically unstable since the numbers quickly become very large. To avoid this problem we instead store the logarithms $\log[\mathcal{Z}]^{\mathrm{e}}_n$. Numpy provides a helpful method `numpy.logaddexp` we can use to iterate to the next step without having to take exponentials. Using this approach equations eq. (6.8) and eq. (6.9) become

$$\log[\mathcal{Z}]^{\mathrm{e}}_n = \texttt{numpy.logaddexp}(\log[\mathcal{Z}]^{\mathrm{e}}_{n-1} + \beta\,{}^{\varepsilon_n}\!/_2, \log[\mathcal{Z}]^{\mathrm{o}}_{n-1} - \beta\,{}^{\varepsilon_n}\!/_2)$$
$$\log[\mathcal{Z}]^{\mathrm{o}}_n = \texttt{numpy.logaddexp}(\log[\mathcal{Z}]^{\mathrm{o}}_{n-1} + \beta\,{}^{\varepsilon_n}\!/_2, \log[\mathcal{Z}]^{\mathrm{e}}_{n-1} - \beta\,{}^{\varepsilon_n}\!/_2)\,.$$

B.1.4. Parallel implementation of parallel tempering

Parallel tempering simulates R replicas of the system evolving at different temperatures. If the evolutions of the replicas were handled sequentially, the runtime would become roughly R times longer. To avoid this additional time cost we can evolve each replica independently and in parallel as different parallel computing processes, with swaps being implemented as communication between these parallel processes. This type of parallel computing is called *message passing* and is available in python through the `mpi4py` class [120]. However, exchanging copies of the system, which in our case are made up of multiple matrices (and actually implemented as a python object), adds a communication cost. We can avoid this by keeping the replicas in place and exchanging the temperatures instead. But this too adds a layer of programming difficulty, in that now we have keep track of which process should communicate with which other since the temperatures are no longer in a fixed order across the parallel processes. I have written a python class `ptmpi` that solves this problem by handling the tracking overhead. (The details of its implementation are given in its github readme.) Additionally, it black boxes all communication between the parallel processes to simplify its use in simulations. The class is independent of the physical system being simulated and could be reused in any other situation to easily implement parallel tempering over parallel processes.

B.1.5. Fast calculation of energy currents using arrays

The sums in the current operator, eq. (5.5), can be rewritten as operations on arrays in order to exploit the massive speed-up of numpy array calculations over python `for` loops [121]. Here the matrix product is denoted with a central dot (`numpy.dot`) and an element-wise product with $*$ (called in numpy with a `*` between arrays). The whole set of energy current expectation values $\mathbf{I} = \langle I_{jk} \rangle$ can be obtained by computing

$$\mathbf{I} = -\frac{1}{8} \left(\mathbf{A} * [\mathbf{C} \cdot \mathbf{A}] + \mathbf{A} * \mathbf{A} * \mathbf{C_{kk}}^T + \mathbf{A} * [\mathbf{A} \cdot \mathbf{C}^T] - \mathbf{A} * \mathbf{A} * \mathbf{C_{kk}} + [\mathbf{A} \cdot \mathbf{A}] * \mathbf{C} \right) ,$$

where $\mathbf{C} = \langle i\, c_j c_k \rangle$ is the matrix of two point correlation functions and $\mathbf{C_{kk}}$ is defined from the diagonal elements of $\mathbf{C}$ by

$$\mathbf{C_{kk}} = \begin{pmatrix} \mathbf{C}_{11} & \mathbf{C}_{22} & \mathbf{C}_{33} & \dots \\ \mathbf{C}_{11} & \mathbf{C}_{22} & \mathbf{C}_{33} & \dots \\ \mathbf{C}_{11} & \mathbf{C}_{22} & \mathbf{C}_{33} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} .$$

This matrix $\mathbf{C_{kk}}$ can be efficiently constructed using numpy's `numpy.tile` method.

B.2. Analysis of Monte Carlo data

Chapter 6 studies the thermal physics of the Kitaev honeycomb spin model through Monte Carlo simulations. This appendix gives a thorough explanation of how the Markov chain time-series obtained from the simulation is processed to give the estimates and errors quoted in the text. This includes accounting for autocorrelation in the time-series, implementing resampling methods to estimate the error on indirectly measured quantities and deciding the equilibration time.

Statistical errors

The Monte Carlo procedure returns a sequence of gauge configurations, $\{u_1, u_2, \ldots u_M\}$. Or, more practically, the values of observables $\mathcal{O}$ (such as fermionic thermal energy, vortex density, etc.) associated with those sectors, $\{O_1, O_2, \ldots O_M\}$. Assuming these values are uncorrelated, the sample mean and (unbiased) sample variance are given respectively by

$$\overline{O} = \mathrm{E}[O_i] = \frac{1}{M}\sum_{i=1}^{M} O_i \,, \tag{B.1}$$

$$s_O^2 = \mathrm{E}\big[\,(O_i - \mathrm{E}[O_i])^2\,\big] = \frac{1}{M-1}\sum_{i=1}^{M}\big(O_i - \overline{O}\big)^2 \,. \tag{B.2}$$

The sample mean $\overline{O}$ estimates the physical expectation value of the observable, eq. (6.18). Error bars on this estimate ΔO are quoted as the standard deviation of $\overline{O}$, computed as

$$(\Delta O)^2 = \mathrm{E}\big[\,(\overline{O} - \mathrm{E}[\overline{O}])^2\,\big] = \frac{s_O^2}{M} \,. \tag{B.3}$$

Autocorrelation times

For Markov chain Monte Carlo simulations the assumption of uncorrelated samples might well be invalid as memory effects can persist in the chains for very long times. These memory effects are quantified by the the *autocorrelation time*, τ, which (roughly) characterises the half-life of the correlations. Autocorrelation effects can cause us to underestimate the error ΔO, since the correlations mean we effectively have fewer independent samples than we believe we do. Or, in a more extreme example, they can bias the estimation

of indirectly measured quantities (such as specific heat), in addition to the error on that estimate [72]. We can gain insight into the autocorrelation between samples by computing the *lagged autocovariance*. For a time series of M samples the lagged autocovariance for observable $\mathcal{O}$ is

$$\gamma_O(t) = \frac{1}{M} \sum_{i=1}^{M-t} \left(O_i - \overline{O}\right) \left(O_{i+t} - \overline{O}\right) ,$$

such that $\gamma_O(0) = s_O^2$. The lagged autocovariance can be used to correct the error ΔO and to estimate the autocorrelation time for the observable τ_O. Additionally, plotting $\gamma_O(t)$ as a function of the lag t can give useful qualitative insights into the decay of the correlations.

The true error, the variance of $\overline{O}$, must include the additional statistical uncertainty introduced by the autocorrelations in the data,

$$\begin{aligned}(\Delta O)^2 = \mathrm{E}\big[\,(\overline{O} - \mathrm{E}[\overline{O}])^2\,\big] &= \frac{1}{M^2} \sum_{i,j=1}^{M} \big(\,\mathrm{E}[O_i O_j] - \mathrm{E}[O_i]\,\mathrm{E}[O_j]\,\big) \\ &= \frac{s_O^2}{M} + \frac{1}{M^2} \sum_{i \neq j} \big(\,\mathrm{E}[O_i O_j] - \mathrm{E}[O_i]\,\mathrm{E}[O_j]\,\big) \end{aligned} \tag{B.4}$$

where the correction term can be computed by summing $\gamma_O(t)$ over all t. Two of the simplest estimators for this correction are the *initial positive sequence estimator* $s^2_{O,\mathrm{p}}$ and the *initial monotone sequence estimator* $s^2_{O,\mathrm{m}}$ [122]. Both aim to sum $\gamma_O(t)$ while accounting for noise that can be problematic at large t. They are constructed from the derived quantities

$$\Gamma_O(t) = \gamma_O(2t) + \gamma_O(2t+1)\,.$$

A theorem states that for stationary, reversible Markov chains these $\Gamma(t)$ should be strictly positive and strictly decreasing. Any violation of these conditions can be interpreted as the onset of noise and used to truncate the sum. The estimator $s^2_{O,\mathrm{p}}$ is then

$$s^2_{O,\mathrm{p}} = s_O^2 + 2 \sum_{t=1}^{2T+1} \gamma_O(t) = -\gamma_O(0) + 2 \sum_{t=0}^{T} \Gamma_O(t) \tag{B.5}$$

where the upper limit T is chosen such that all the terms summed are positive, $\Gamma_O(t) > 0$ for $t = 0, \ldots T$. Estimator $s^2_{O,\mathrm{m}}$ is defined similarly, with the upper limit decreased further to the index of the minimum preceding value of $\Gamma_O(t)$ so that the sequence of

terms summed $\Gamma_O(0)$, $\Gamma_O(1)$, ... is both always positive and always decreasing. Dividing these estimators by M estimates the true error,

$$(\Delta O)^2 = \frac{s^2_{O,\mathrm{m}}}{M}\,. \tag{B.6}$$

Assuming the autocorrelations decay exponentially, $\gamma_O(t) \sim e^{-t/\tau^*_O}$, the exponential decay rate can also be computed from the integral of $\gamma_O(t)$ as

$$\int_0^\infty \mathrm{d}t\, \frac{\gamma_O(t)}{\gamma_O(0)} = \int_0^\infty \mathrm{d}t\, e^{-t/\tau^*_O} = \tau^*_O\,.$$

Comparing this to eq. (B.5), where the sum over t is viewed as the discretised integral, we see that dividing the improved variance estimator $s^2_{O,\mathrm{m}}$ by the raw sample variance s^2_O yields a related estimate of the autocorrelation time (called the *integrated autocorrelation time*) [72, 123] defined as

$$\overline{\tau}_O \equiv \frac{1}{2}\frac{s^2_{O,\mathrm{m}}}{s^2_O} = \tau^*_O - \frac{1}{2}\,, \tag{B.7}$$

This small difference between the exponential decay rate τ^*_O and $\overline{\tau}_O$ is typically neglected. In the following discussions the integrated autocorrelation time $\overline{\tau}_O$ is the value quoted as the autocorrelation time τ_O of the Markov chains. Combining the expressions for ΔO we can also define the effective number of samples,

$$M_{\mathrm{eff}} = \frac{M}{2\,\overline{\tau}_O}\,, \tag{B.8}$$

such that eq. (B.6) instead reads $(\Delta O)^2 = s^2_O/M_{\mathrm{eff}}$.

Bootstrap resampling

Assigning an error to estimates of indirectly measured quantities, such as distributional parameters like the variance, is more complicated than for quantities that are directly measured by the simulation. The easiest approaches are resampling methods including *blocking*, *jackknife* and *bootstrap* [124, 125]. The bootstrap method is among the most robust as well as (for short enough autocorrelation times) the most efficient.

From the gauge configuration time series, $\{u_1, u_2, \ldots\, u_M\}$, we select M random elements

allowing repeats, e.g. $\mathcal{B}_1 = \{1, 1, 3, 4, 4, 4, 4, 8 \ldots\}$. The quantity is computed for this sample, for example the specific heat is evaluated from the sample variance

$$C_1 = \frac{1}{T^2}\left(\frac{1}{M-1}\sum_{i\in\mathcal{B}_1}(E_i - \overline{E})\right) \quad \text{where} \quad \overline{E} = \frac{1}{M}\sum_{i\in\mathcal{B}_1} E_i\,. \tag{B.9}$$

The time series is then resampled repeatedly, each time with a new set of random elements $\mathcal{B}_2$, $\mathcal{B}_3$, ... to compute a series of estimates $\{C_1, C_2, C_3, \ldots\}$. For a large enough number of resamples R, typically between 100 and 1000, the mean of these estimates gives a good estimator of the specific heat with the error ΔC given by their standard deviation,

$$\overline{C} = \frac{1}{R}\sum_{i=1}^{R} C_i \quad \text{with} \quad (\Delta C)^2 = \overline{C^2} - \overline{C}^2\,. \tag{B.10}$$

This statistical witchcraft has the advantage that, since the resampling is already carried out with repeats, it is not effected by autocorrelation in the data [125]. We may use the full Monte Carlo time series to compute $\overline{C}$ and ΔC.

Equilibration

The Markov chain must be allowed to equilibrate before we collect data. Correlations between the initial state and future states can behave differently to autocorrelations present along all the data, meaning the equilibration time can be much longer than τ [126].

Waiting for equilibration is important even if we could pick an initial state that is guaranteed to be part of the thermal distribution at all temperatures, such as the true ground state. This is because a histogram of (for example) the energies of states the Markov chain visits is attempting to reproduce the physical probability distribution of energies. If the initial configuration is in the tail of the distribution then the Markov chain histogram will be skewed. Using states of the Markov chain prior to equilibration to estimate properties of the physical distribution, such as its mean (expected energy) and variance (specific heat), will then bias these estimates. Thus adding bias errors in addition to the statistical estimation errors.

However, since the statistical error decreases with increasing numbers of data points we do not want to unnecessarily throw away too much of the start of the Markov chain. A

good way to resolve this trade-off is to fix the equilibration time t_0 as the balance point between minimising bias error and minimising statistical error [127]. The effective number of samples M_{eff} as a function of t_0 is defined as

$$M_{\text{eff}}(t_0) = \frac{M - t_0}{2\,\overline{\tau}_o(t_0)} \tag{B.11}$$

where $\overline{\tau}_o(t_0)$ is the integrated equilibration time computed for the time series after t_0, $\{O_{t_0+1}, O_{t_0+2}, \ldots O_M\}$. A balance between bias and statistical error is said to be achieved by maximising $M_{\text{eff}}(t_0)$ with respect to t_0. Each observables $\mathcal{O}$ equilibrates at a different rate, and true equilibration is only achieved when all observables have reached equilibrium.

For a simulation using parallel tempering a simpler estimate of the equilibration time is the mean first passage time [128, 129]. This is the average time it takes a replica initially at one end of the temperature set to diffuse to the other for the first time. Rather than taking a different value for each observable, this estimate of the equilibration time is a property of the simulation as a whole.

Summary

The simulation returns a sequence of values for each of the observables we are interested in, $\{O_1, O_2, \ldots O_M\}$. The equilibration time t_0 of the simulation run is estimated from the first passage times of the parallel tempering replicas. We evaluate t_0 as the average of the time the replica initially at lowest temperature first reaches highest temperature and the time the initial highest temperature replica first reaches lowest temperature. The first t_0 values of the time series are discarded.

Mean values of measured observables are then estimated as the sample mean, eq. (B.1), of the remaining time series. The error on this estimate is given by eq. (B.6), where $s^2_{O,\text{m}}$ is again computed using the remaining time series and M is reduced to the number of elements after discarding the first t_0. Indirectly measured properties such as the variance of physical observables are estimated, along with an error, using bootstrap resampling, eq. (B.10), with 1000 resamples of the remaining time series.

C.

Dynamics of Kitaev's Honeycomb

Some of the most important physics of a model is its response to external perturbations. These responses can be probed in experiments to compare real materials to the theory. In this way we gain insight into whether the simplified model captures the real behaviours or if other competing mechanisms, neglected by the model, cannot be ignored.

For the Kitaev Honeycomb the relevant responses are heat currents reacting to a thermal gradient. Though temperature is a statistical effect, thermal gradients can be simulated as a Hamiltonian perturbation and analysed within a conventional response theory setting [130]. Let the initial fermionic Hamiltonian be H_0 with thermal state ρ_0 and consider a time-dependent perturbation H'. This perturbation couples the energy density $h(\mathbf{r})$ to a potential field $\phi(\mathbf{r})$. If the time-dependence of H' is $u(t)$, then the total Hamiltonian of the system is

$$H = H_0 + H'(t) = H_0 + u(t) \int \mathrm{d}\mathbf{r}\, h(\mathbf{r})\phi(\mathbf{r})\,. \tag{C.1}$$

The perturbation causes energy currents to flow in the system, whose expectation values are the vector $\langle \mathbf{J}(t) \rangle$. Specialising to potentials that vary linearly between the two ends of a cylindrical system (at $y = 0$ and $y = L$), the spatial dependence of the perturbation can be written in vector component notation $(\nabla\phi)_\gamma = {}^{\Delta\phi}/_{L}\, \delta_{\gamma y}$. For small enough perturbation the response in the system is *linear*. In this limit, the total current that flows in direction α is given by

$$\begin{aligned} \langle J_\alpha(t) \rangle &= \int_{-\infty}^{\infty} \mathrm{d}t' \sum_\gamma \sigma_{\alpha\gamma}(t-t')\, u(t')\, (\nabla\phi)_\gamma \\ &= \int_{-\infty}^{\infty} \mathrm{d}t'\, \sigma_{\alpha y}(t-t')\, u(t')\, {}^{\Delta\phi}/_{L}\,, \end{aligned} \tag{C.2}$$

where it was shown by Luttinger [130] that the conductivity tensor $\sigma_{\alpha\gamma}$ can be written in the form of a Kubo formula as

$$\sigma_{\alpha y}(t) = \theta(t) \int_0^\beta \mathrm{d}\beta'\, \mathrm{tr}\left(\rho_0\, J_y(-i\beta')\, J_\alpha(t)\right)\,. \tag{C.3}$$

C.1. Kubo Formula

Under linear response, the system's reaction is fully characterised by its Green's functions. Importantly, these Green's function can be computed using the unperturbed wavefunctions through the Kubo formula. Rewriting eq. (C.3) in a slightly different form, the conductivity is described by the tensor

$$\sigma_{\alpha\gamma}(t) \;=\; \theta(t) \int_0^\beta \mathrm{d}\beta' \, \langle J_\gamma(-t - i\beta') \, J_\alpha \rangle \,, \tag{C.4}$$

where the expectation values are computed with respect to the unperturbed thermal state at temperature β. Neglecting the intergral over imaginary time, the response is determined by the two-time current-current correlations $\langle J_\gamma(t') \, J_\alpha \rangle$. In this section we consider how these correlations can be computed using the real space energy current operators defined in chapter 5.

C.1.1. Current-current correlators

Revisiting the construction of the Honeycomb model from chapter 4, the two-point Majorana correlation functions are given by

$$\begin{aligned}
i\, c_j c_k &= 2i \sum_{p,q} \Big([\underline{u}_p]_j \, f_p^\dagger + [\underline{u}_p]_j^* \, f_p \Big) \Big([\underline{u}_q]_k \, f_q^\dagger + [\underline{u}_q]_k^* \, f_q \Big) \\
&= 2i \sum_{p,q} \Big([\underline{u}_p]_j [\underline{u}_q]_k \, f_p^\dagger f_q^\dagger + [\underline{u}_p]_j [\underline{u}_q]_k^* \, f_p^\dagger f_q + [\underline{u}_p]_j^* [\underline{u}_q]_k \, f_p f_q^\dagger + [\underline{u}_p]_j^* [\underline{u}_q]_k^* \, f_p f_q \Big) \\
&= 2i \sum_{p,q} \Big([\underline{u}_p \, (\underline{u}_q^*)^\dagger]_{jk} \, f_p^\dagger f_q^\dagger + [\underline{u}_p (\underline{u}_q)^\dagger]_{jk} \, f_p^\dagger f_q + [\underline{u}_p (\underline{u}_q)^\dagger]_{jk}^* \, f_p f_q^\dagger + [\underline{u}_p (\underline{u}_q^*)^\dagger]_{jk}^* \, f_p f_q \Big) \,.
\end{aligned}$$

The weightings attached to each set of operators, for example $[\underline{u}_p \, (\underline{u}_q^*)^\dagger]_{jk}$, are matrices associated with the lattice wavefunctions. The operators $i\, c_j c_k$ are Hermitian observables and this puts conditions on the matrices. By comparing the expression above to that for

$(i\,c_j c_k)^\dagger$ we find they must satisfy the relations:

$$\left[\underline{u}_q\,(\underline{u}_p)^\dagger\right]_{ij} = -\left[\underline{u}_p\,(\underline{u}_q)^\dagger\right]^*_{ij} \tag{C.5}$$

$$\left[\underline{u}_q\,(\underline{u}_p)^\dagger\right]^*_{ij} = -\left[\underline{u}_p\,(\underline{u}_q)^\dagger\right]_{ij} \tag{C.6}$$

$$\left[\underline{u}_q\,(\underline{u}^*_p)^\dagger\right]_{ij} = -\left[\underline{u}_p\,(\underline{u}^*_q)^\dagger\right]_{ij} \tag{C.7}$$

$$\left[\underline{u}_q\,(\underline{u}^*_p)^\dagger\right]^*_{ij} = -\left[\underline{u}_p\,(\underline{u}^*_q)^\dagger\right]^*_{ij}\,. \tag{C.8}$$

(Numerically, we must ensure these are obeyed. This can be done in the way described in note B.1.1.) These relations will be used in the following to simplify expressions.

The real space current operator derived in chapter 5 from the Heisenburg equation of local energy densities is written

$$I_{jk} \;=\; \frac{i}{8}\left(\sum_{m\neq k} A_{kj}A_{jm}\,c_k c_m + \sum_{l\neq j} A_{kl}A_{jk}\,c_l c_j + \sum_{l\neq j} A_{kl}A_{jl}\,c_j c_k\right)\,.$$

This can be rewritten in the energy eigenbasis using the expression for $i\,c_j c_k$ above. However, to hide some of the complexity, first define

$$\mathcal{Q}_{jk}[X_{ab}] \equiv -\frac{i}{8}\left(\sum_{m\neq k} A_{kj}A_{jm}\,X_{km} + \sum_{l\neq j} A_{kl}A_{jk}\,X_{lj} + \sum_{l\neq j} A_{kl}A_{jl}\,X_{jk}\right)\,. \tag{C.9}$$

The current operators in the energy eigenbasis then take the form

$$\begin{aligned} I_{jk} = \sum_{p,q}\Big(\; & \mathcal{Q}_{jk}\Big[\underline{u}_p\,(\underline{u}^*_q)^\dagger\Big]\; f^\dagger_p f^\dagger_q + \mathcal{Q}_{jk}\Big[\underline{u}_p(\underline{u}_q)^\dagger\Big]\; f^\dagger_p f_q \\ & + \mathcal{Q}_{jk}\Big[\Big(\underline{u}_p(\underline{u}_q)^\dagger\Big)^*\Big]\; f_p f^\dagger_q + \mathcal{Q}_{jk}\Big[\Big(\underline{u}_p(\underline{u}^*_q)^\dagger\Big)^*\Big]\; f_p f_q\Big)\,. \end{aligned} \tag{C.10}$$

However, the current-current correlation functions we want to evaluate $\langle J_\alpha(t)\,J_\beta\rangle$ are between the *total* currents J. Total currents are defined as a sum of I_{jk} over all links of one type α,

$$\begin{aligned} J_\alpha \equiv \sum_{(j,k)\,:\,\alpha} I_{jk} \;= \sum_{p,q}\Big(\; & \mathcal{Q}_\alpha\Big[\underline{u}_p\,(\underline{u}^*_q)^\dagger\Big]\; f^\dagger_p f^\dagger_q + \mathcal{Q}_\alpha\Big[\underline{u}_p(\underline{u}_q)^\dagger\Big]\; f^\dagger_p f_q \\ & + \mathcal{Q}_\alpha\Big[\Big(\underline{u}_p(\underline{u}_q)^\dagger\Big)^*\Big]\; f_p f^\dagger_q + \mathcal{Q}_\alpha\Big[\Big(\underline{u}_p(\underline{u}^*_q)^\dagger\Big)^*\Big]\; f_p f_q\Big)\,, \end{aligned}$$

where $\mathcal{Q}_\alpha[X] = \sum_{(j,k)\,:\,\alpha} Q_{jk}[X]$. Further, the correlation function is between the currents at two different times so we need to evolve the operator J in time. Since we are now in the energy eigenbasis this is easy to do,

$$\begin{aligned} J_\alpha(t) &= e^{iHt}\, J_\alpha\, e^{-iHt} \\ &= \sum_{p,q} \Big(\mathcal{Q}_\alpha\Big[\underline{u}_p\,(\underline{u}_q^*)^\dagger\Big]\, e^{i(\varepsilon_p+\varepsilon_q)t/2}\, f_p^\dagger f_q^\dagger + \mathcal{Q}_\alpha\Big[\underline{u}_p(\underline{u}_q)^\dagger\Big]\, e^{i(\varepsilon_p-\varepsilon_q)t/2}\, f_p^\dagger f_q \\ &\qquad + \mathcal{Q}_\alpha\Big[\Big(\underline{u}_p(\underline{u}_q)^\dagger\Big)^*\Big]\, e^{-i(\varepsilon_p+\varepsilon_q)t/2}\, f_p f_q^\dagger + \mathcal{Q}_\alpha\Big[\Big(\underline{u}_p(\underline{u}_q^*)^\dagger\Big)^*\Big]\, e^{-i(\varepsilon_p+\varepsilon_q)t/2}\, f_p f_q \Big)\ . \end{aligned}$$

Expanding the operator expression for $J_\alpha(t)\, J_\beta$ using these formulae gives the following Frankenstein monster of an equation

$$\begin{aligned} J_\alpha(t)\, J_\beta = \sum_{p,q}\sum_{r,s} \Big(& \mathcal{Q}_\alpha\Big[\underline{u}_p\,(\underline{u}_q^*)^\dagger\Big]\, \mathcal{Q}_\beta\Big[\Big(\underline{u}_r(\underline{u}_s^*)^\dagger\Big)^*\Big]\ f_p^\dagger f_q^\dagger f_r f_s\, e^{i(\varepsilon_p+\varepsilon_q)t/2} \\ & + \mathcal{Q}_\alpha\Big[\underline{u}_p(\underline{u}_q)^\dagger\Big]\, \mathcal{Q}_\beta\Big[\underline{u}_r(\underline{u}_s)^\dagger\Big]\ f_p^\dagger f_q\, f_r^\dagger f_s\, e^{i(\varepsilon_p-\varepsilon_q)t/2} \\ & + \mathcal{Q}_\alpha\Big[\underline{u}_p(\underline{u}_q)^\dagger\Big]\, \mathcal{Q}_\beta\Big[\Big(\underline{u}_r(\underline{u}_s)^\dagger\Big)^*\Big]\ f_p^\dagger f_q\, f_r f_s^\dagger\, e^{i(\varepsilon_p-\varepsilon_q)t/2} \\ & + \mathcal{Q}_\alpha\Big[\Big(\underline{u}_p(\underline{u}_q)^\dagger\Big)^*\Big]\, \mathcal{Q}_\beta\Big[\underline{u}_r(\underline{u}_s)^\dagger\Big]\ f_p f_q^\dagger\, f_r^\dagger f_s\, e^{-i(\varepsilon_p-\varepsilon_q)t/2} \\ & + \mathcal{Q}_\alpha\Big[\Big(\underline{u}_p(\underline{u}_q)^\dagger\Big)^*\Big]\, \mathcal{Q}_\beta\Big[\Big(\underline{u}_r(\underline{u}_s)^\dagger\Big)^*\Big]\ f_p f_q^\dagger\, f_r f_s^\dagger\, e^{-i(\varepsilon_p-\varepsilon_q)t/2} \\ & + \mathcal{Q}_\alpha\Big[\Big(\underline{u}_p(\underline{u}_q^*)^\dagger\Big)^*\Big]\, \mathcal{Q}_\beta\Big[\Big(\underline{u}_r(\underline{u}_s^*)^\dagger\Big)^*\Big]\ f_p f_q\, f_r^\dagger f_s^\dagger\, e^{-i(\varepsilon_p+\varepsilon_q)t/2}\Big)\ . \end{aligned} \tag{C.11}$$

The terms written are only those that include two annihilation and two creation operators, since any other combinations of four operators will have zero expectation value.

Evaluating the Kubo formula requires taking expectation values of eq. (C.11). In order to do this we must compute the matrix functions $\mathcal{Q}_\alpha[X]$, which we will refer to as *lattice factors*, and take the expectation values of the four fermion operators. Since there are L^2 unique wavefunctions $\underline{u}_p$, it seems as if there are $4L^4$ lattice factors $\mathcal{Q}_\alpha$ for each direction α (counting all combinations of conjugations). However, not all of these lattice factors are independent. From their definition, eq. (C.9), we see that $\mathcal{Q}_\alpha[X]^* = -\mathcal{Q}_\alpha[X^*]$. Further, since the current operator, eq. (C.10), is real the lattice factors must themselves be real

$\mathcal{Q}_\alpha[X]^* = \mathcal{Q}_\alpha[X]$. This gives the relation

$$\mathcal{Q}_\alpha[X^*_{ab}] = -\mathcal{Q}_\alpha[X_{ab}]\,. \tag{C.12}$$

Additionally, conditions eq. (C.5)-(C.8) limit the number of independent $[\underline{u}_p\,(\underline{u}_q)^\dagger]$ and $[\underline{u}_p\,(\underline{u}^*_q)^\dagger]$ matrices to $L^2(L^2+1)/2$ each. Hence the total number of lattice factors we must calculate is $L^2(L^2+1)$.

Next consider the thermal expectation values of the products of four fermion operators. A normal ordered set of four operators can be seen to give

$$\begin{aligned}
\langle a_1^\dagger a_2^\dagger\, a_3 a_4\rangle &= \delta_{13}\delta_{24}\,\langle a_1^\dagger a_2^\dagger\, a_1 a_2\rangle\,(1-\delta_{12}) \\
&\quad + \delta_{14}\delta_{23}\,\langle a_1^\dagger a_2^\dagger\, a_2 a_1\rangle\,(1-\delta_{12}) \\
&= \delta_{13}\delta_{24}\,(-1)\,\langle a_1^\dagger a_1\, a_2^\dagger a_2\rangle\,(1-\delta_{12}) \\
&\quad + \delta_{14}\delta_{23}\,\langle a_1^\dagger a_1\, a_2^\dagger a_2\rangle\,(1-\delta_{12}) \\
&= (1-\delta_{12})\,\langle a_1^\dagger a_1\, a_2^\dagger a_2\rangle\,(\delta_{13}\delta_{24}-\delta_{14}\delta_{23})\,,
\end{aligned} \tag{C.13}$$

where, due to the projection requirement of the Kitaev fermions, it is not necessarily the case that $\langle a_1^\dagger a_1\, a_2^\dagger a_2\rangle = \langle a_1^\dagger a_1\rangle\langle a_2^\dagger a_2\rangle$. Rearranging each of the six terms in eq. (C.11) into normal order and applying eq. (C.13) we find that

$$\begin{aligned}
\langle f_p^\dagger f_q^\dagger f_r f_s\rangle &= \langle f_p^\dagger f_p\, f_q^\dagger f_q\rangle\,(\delta_{ps}\delta_{qr}-\delta_{pr}\delta_{qs})\,(1-\delta_{pq}) \\
\langle f_p^\dagger f_q f_r^\dagger f_s\rangle &= \Big(\delta_{rq}\delta_{ps}\,\langle f_p^\dagger f_p\, f_r f_r^\dagger\rangle + \delta_{pq}\delta_{rs}\,\langle f_p^\dagger f_p\, f_r^\dagger f_r\rangle\Big)\,(1-\delta_{pr}) \\
\langle f_p^\dagger f_q f_r f_s^\dagger\rangle &= \langle f_p^\dagger f_p\, f_s f_s^\dagger\rangle\,(\delta_{pq}\delta_{rs}-\delta_{pr}\delta_{qs})\,(1-\delta_{ps}) \\
\langle f_p f_q^\dagger f_r^\dagger f_s\rangle &= \Big(\delta_{pq}\delta_{rs}\,\langle f_q f_q^\dagger\, f_r^\dagger f_r\rangle - \delta_{qs}\delta_{pr}\,\langle f_q^\dagger f_q\, f_r f_r^\dagger\rangle\Big)\,(1-\delta_{qr}) \\
\langle f_p f_q^\dagger f_r f_s^\dagger\rangle &= \Big(\delta_{pq}\delta_{rs}\,\langle f_q f_q^\dagger\, f_s f_s^\dagger\rangle + \delta_{ps}\delta_{qr}\,\langle f_q^\dagger f_q\, f_s f_s^\dagger\rangle\Big)\,(1-\delta_{qs}) \\
\langle f_p f_q f_r^\dagger f_s^\dagger\rangle &= \langle f_r f_r^\dagger\, f_s f_s^\dagger\rangle\,(\delta_{rq}\delta_{ps}-\delta_{pr}\delta_{qs})\,(1-\delta_{rs})\,.
\end{aligned}$$

Using these expressions, combined with eq. (C.12), the expectation values of the current-

current correlations can be written

$$\begin{aligned}\langle J_\alpha(t)\, J_\beta\rangle = \sum_{p\neq q} \mathcal{Q}_\alpha\Big[\underline{u}_p\,(\underline{u}_p)^\dagger\Big]\,\mathcal{Q}_\beta\Big[\underline{u}_q(\underline{u}_q)^\dagger\Big]\,\Big(\langle f_p^\dagger f_p\, f_q^\dagger f_q\rangle + \langle f_p f_p^\dagger\, f_q f_q^\dagger\rangle \\ -\langle f_p^\dagger f_p\, f_q f_q^\dagger\rangle - \langle f_p f_p^\dagger\, f_q^\dagger f_q\rangle\Big) \\ + 2\sum_{p\neq q}\Big(\langle f_p f_p^\dagger\, f_q^\dagger f_q\rangle\, e^{i(\varepsilon_p+\varepsilon_q)t/2}\,\mathcal{Q}_\alpha\Big[\underline{u}_p\,(\underline{u}_q^*)^\dagger\Big]\,\mathcal{Q}_\beta\Big[\underline{u}_p(\underline{u}_q^*)^\dagger\Big] \\ + \langle f_p^\dagger f_p\, f_q f_q^\dagger\rangle\, e^{i(\varepsilon_p-\varepsilon_q)t/2}\,\mathcal{Q}_\alpha\Big[\underline{u}_p\,(\underline{u}_q)^\dagger\Big]\,\mathcal{Q}_\beta\Big[\underline{u}_p(\underline{u}_q)^\dagger\Big] \\ + \langle f_p f_p^\dagger\, f_q^\dagger f_q\rangle\, e^{-i(\varepsilon_p-\varepsilon_q)t/2}\,\mathcal{Q}_\alpha\Big[\underline{u}_p\,(\underline{u}_q)^\dagger\Big]\,\mathcal{Q}_\beta\Big[\underline{u}_p(\underline{u}_q)^\dagger\Big] \\ + \langle f_p f_p^\dagger\, f_q f_q^\dagger\rangle\, e^{-i(\varepsilon_p+\varepsilon_q)t/2}\,\mathcal{Q}_\alpha\Big[\underline{u}_p(\underline{u}_q^*)^\dagger\Big]\,\mathcal{Q}_\beta\Big[\underline{u}_p(\underline{u}_q^*)^\dagger\Big]\Big)\,. \end{aligned} \tag{C.14}$$

Integrating this expression over imaginary time, as in eq. (C.4), then gives the thermal conductivities.

C.1.2. Frequency response

The response function in frequency space $\sigma_{\alpha y}(\omega)$, the Fourier transform of $\sigma_{\alpha y}(t)$, is an even more interesting quantity. In particular, it can tell us if the system is an insulator or a conductor. Numerically, since we can only compute the time-series $\sigma_{\alpha y}(t)$ up to finite $t_{\max}$, it is better to evaluate the Laplace transform in the limit $\delta \to 0$,

$$\sigma_{\alpha y}(\omega) = \lim_{\delta\to 0}\int_0^{t_{\max}} \mathrm{d}t\, e^{-(\delta+i\omega)t}\,\sigma_{\alpha y}(t)\,,$$

where as $\delta \to 0$ we increase $t_{\max}$, rather than attempt the Fourier transform directly.

In a clean model system with no impurities a conductor is characterised by a non-zero Drude weight in the zero temperature conductivity [131]. The Drude weight, D, manifests as a delta function at $\omega = 0$ in the frequency response

$$\sigma_{\alpha y}(\omega; T=0) = D_\alpha\,\delta(\omega) + \sigma_{\alpha y}^{\mathrm{reg}}(\omega; T=0)\,. \tag{C.15}$$

This produces dissipationless currents in the impurity free system at zero temperature. An insulator, in contrast, has vanishing Drude weight. We can extract the Drude weight

from the time-series $\sigma_{\alpha y}(t)$ by computing

$$\overline{D}_\alpha = \int_{t_1}^{t_2} \mathrm{d}t \, \frac{\sigma_{\alpha y}(t)}{t_2 - t_1} \tag{C.16}$$

where t_1 and t_2 are chosen much longer than any timescale of the system [132].

C.2. Proposal: Quantum evolution after a quench

The conductivity tensor $\sigma_{\alpha y}(t)$ is the Green's function $g_\alpha(t)$ for the currents J_α, eq. (C.2). Physically this means it is the current that flows in response to an impulse shock to the system $u(t) = \delta(t)$. This section proposes an alternative approach to calculating the conductivity by directly computing the response currents $\langle J_\alpha(t) \rangle$ from the time evolution of the wavefunctions under a quenched perturbation. This calculation can be carried out for a particular perturbation and, assuming the response is linear, the Green's function can be extracted from Laplace transforms of the currents. The proposal involves simulating the full quantum evolution of the perturbed system. Choosing a simple quenched perturbation that switches on at $t = 0$ (a step function) makes this easy to calculate. Besides being an alternative way to compute linear response, this approach might be interesting for probing non-linear aspects of the response.

Let the time dependence of the perturbation be a step-function $u(t) = \theta(t)$. In this case, the total Hamiltonian H, given by eq. (C.1), is time-independent even after the perturbation has been switched on. By evolving the initial state ρ_0 under H the current at times $t > 0$ (normalised by system size and strength of the perturbation) is given by

$$\begin{aligned} I_\alpha(t;\theta) &\equiv \frac{1}{L^2} \left(\frac{\Delta\phi}{L} \right)^{-1} \mathrm{tr} \left(e^{-iHt} \, \rho_0 \, e^{iHt} \, J_\alpha \right) \\ &= \frac{1}{L^2} \left(\frac{\Delta\phi}{L} \right)^{-1} \sum_{n,m} e^{-i(\varepsilon_n - \varepsilon_m)t} \, \langle n | \, \rho_0 \, | m \rangle \, \langle m | \, J_\alpha \, | n \rangle \end{aligned} \tag{C.17}$$

where $|n\rangle, |m\rangle$ and $\varepsilon_n, \varepsilon_m$ are the eigenstates and eigenenergies of H. We can compute this just by diagonalising H_0 and H once, and (in principle) find the current $I_\alpha(t;\theta)$ at all times. For an arbitrary input function $u(t)$ and response current $I_\alpha(t;u)$ we can find the

Green's function $g_\alpha(t)$ by evaluating

$$g_\alpha(t) = \mathcal{L}^{-1}\left[\frac{\mathcal{L}[I_\alpha(t;u)]}{\mathcal{L}[u]}\right] \tag{C.18}$$

where $\mathcal{L}$ is the Laplace transform [133], which for complex s is given by

$$\mathcal{L}\big[u\big](s) = \int_0^\infty \mathrm{d}t\, e^{-st}\, u(t)\,.$$

By our choice of $u(t) = \theta(t)$ eq. (C.18) is easy to evaluate. In this case $\mathcal{L}[\theta] = {}^1\!/_s$. Using an identity for Laplace transforms the Green's function is given by

$$\begin{aligned} g_\alpha(t) &= \big(I_\alpha(t;\theta) \,*\, \mathcal{L}^{-1}[s]\big)(t) \\ &= \int_\infty^\infty \mathrm{d}t'\, \delta'(t')\, I_\alpha(t-t';\theta) \\ &= \int_\infty^\infty \mathrm{d}t'\, \delta(t')\, \dot{I}_\alpha(t-t';\theta) = \dot{I}_\alpha(t;\theta) \end{aligned} \tag{C.19}$$

where $*$ is a convolution, δ' is the distributional-derivative of the delta function (which is the inverse Laplace transform of s) and $\dot{I}_\alpha$ is a time-derivative. The Green's function is causal, which we can explicitly denote by including a step-function $g_\alpha(t) = \theta(t)\dot{I}_\alpha(t;\theta)$. For completeness, the current response for an arbitrary $u(t)$ is then

$$I_\alpha(t;u) = \int_{-\infty}^\infty \mathrm{d}t'\, g_\alpha(t-t')\, u(t')\,.$$

To reiterate, the Green's function given in eq. (C.19) is the same quantity as the conductivity tensor previously defined $\sigma_{\alpha y}(t) \equiv g_\alpha(t)$. With the caveat that, since we are computing the currents from the full quantum evolution, the response is not guaranteed to be purely linear. In some cases the nonlinear contributions this approach can highlight might be interesting in themselves. Otherwise, we must separately ensure we are in the linear regime in order to apply eq. (C.18) and be computing the same conductivity that we would find from the Kubo formula.

Frequency dependence

The frequency dependence of the conductivity $\sigma_{\alpha y}(\omega)$ can be obtained without needing to carry out the differentiation of $I_\alpha(t;\theta)$. Consider the Laplace transform of $\sigma_{\alpha y}(t)$ given by

$$\mathcal{L}\big[\sigma_{\alpha y}(t)\big](s=\delta+i\omega)=\int_0^\infty \mathrm{d}t\, e^{-st}\,\dot{I}_\alpha(t;\theta)\,,$$

which as $\delta\to 0$ becomes the Fourier transform. Using an identity for the Laplace transform of a derivative this simplifies to

$$\sigma_{\alpha y}(\omega)=\lim_{\delta\to 0}\, s\int_0^\infty \mathrm{d}t\, e^{-st}\, I_\alpha(t;\theta)\,.$$

D.C. conductivity

Special care needs to be taken in computing the d.c. component, $\omega=0$, of $\sigma_{\alpha y}(\omega)$ because of the behaviour of $\mathcal{L}[I_\alpha(t;\theta)]$ close to $s=0$. In the case $\sigma_{\alpha y}(0)\neq 0$ the current $I_\alpha(t;\theta)$ is not square-integrable because it has a non-zero mean value. We can separate out this mean value and write

$$I_\alpha(t;\theta)=D\,\theta(t)+\tilde{I}_\alpha(t;\theta)$$

where $\tilde{I}_\alpha(t;\theta)$ has mean value zero. The Laplace transform at zero-frequency then has the form

$$\mathcal{L}\big[I_\alpha(t;\theta)\big](s=\delta+i0)=\frac{D}{\delta}+\mathcal{L}\big[\tilde{I}_\alpha(t;\theta)\big]\,.$$

Comparing this to section C.2 we see this gives the desired result

$$\sigma_{\alpha y}(0)=\lim_{\delta\to 0}\,\delta\,\Big(\,\frac{D}{\delta}+\mathcal{L}\big[\tilde{I}_\alpha(t;\theta)\big]\,\Big)=D\,.$$

However, if we have failed to make δ small enough when computing $\mathcal{L}[I_\alpha(t;\theta)]$ and are not in a regime where there is $\sim {}^1\!/_\delta$ scaling the factors will not cancel correctly and we will find the wrong value for $\sigma_{\alpha y}(0)$. To avoid this problem we compute $\mathcal{L}[I_\alpha(t;\theta)]$ for different small values of δ and check that we see $\sim {}^1\!/_\delta$ behaviour.

C.2.1. Limited time pulse

Instead of a step-function we now consider an interaction that turns on at $t = 0$ and turns off again at $t = T$. That is a time-dependence $u(t) = {}^1\!/_T\,\theta(t)\theta(T-t)$ we call this $P_{0,T}(t)$, which has Laplace transform

$$\mathcal{L}[P_{0,T}] = \frac{1}{T}\frac{1-e^{-sT}}{s}\,.$$

Following the prescription of eq. (C.18) and eq. (C.19), in order to compute the conductivity we need to find the Inverse Laplace transform of ${}^1\!/_{\mathcal{L}[P_{0,T}]}$. We do this explicitly by computing the integral

$$\mathcal{L}^{-1}\left[\frac{1}{\mathcal{L}[P_{0,T}]}\right] = \frac{1}{2\pi i}\lim_{R\to\infty}\int_{-iR}^{iR} e^{st}\frac{T\,s}{1-e^{-sT}}\,\mathrm{d}s\,.$$

Since $t > 0$ we close the contour in the left-half-plane. By residue theorem we can then evaluate the integral by finding the residues at the poles of the integrand. The poles will come from the ${}^1\!/_{1-e^{-sT}}$ part of the integrand, to locate them we rewrite it as

$$\frac{1}{1-e^{-sT}} = \frac{1}{2e^{-sT/2}}\,\mathrm{csch}\left(\frac{sT}{2}\right)\,.$$

$\mathrm{csch}(z)$ has an infinite number of simple poles on the imaginary line at $z = \pi k i$ for $k \in \mathbb{Z}$ with residues $(-1)^k$, meaning our integrand has poles at $s_k = {}^{2\pi i k}\!/_T$. Combining this with the rest of the integrand we find

$$\begin{aligned}
\mathcal{L}^{-1}\left[\frac{1}{\mathcal{L}[P_{0,T}]}\right] &= \sum_k \lim_{s\to s_k}(s-s_k)\frac{1}{2}s\,e^{s(t+T/2)}\mathrm{csch}\left(\frac{sT}{2}\right)\\
&= \sum_{k=-\infty}^{\infty}\frac{1}{2}s_k\,e^{s_k(t+T/2)}(-1)^k\\
&= \sum_{k=-\infty}^{\infty}(-1)^k\frac{1}{2}\frac{2\pi i k}{T}\exp\left(\frac{2\pi i k}{T}(t+{}^T\!/_2)\right)\\
&= \frac{1}{2}\sum_{k=-\infty}^{\infty}(-1)^k(e^{i\pi})^k\frac{2\pi i k}{T}\exp\left(\frac{2\pi i k}{T}t\right)\\
&= \frac{1}{2}\sum_{k=-\infty}^{\infty}\frac{2\pi i k}{T}\exp\left(\frac{2\pi i k}{T}t\right)\\
&= \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{D}(t)
\end{aligned}$$

where $\mathcal{D}(t)$ is the Dirac comb. Inserting this into the expression for $g_\alpha(t)$, eq. (C.18), gives

$$\begin{aligned} g_\alpha(t) &= \int_{-\infty}^{\infty} \mathrm{d}t'\, I_\alpha(t-t'; P_{0,T})\, \frac{1}{2}\, \frac{\mathrm{d}}{\mathrm{d}t}\mathcal{D}(t) \\ &= \frac{1}{2}\int_{-\infty}^{\infty} \mathrm{d}t'\, \dot{I}_\alpha(t-t'; P_{0,T})\, \mathcal{D}(t) \\ &= \frac{1}{2}\sum_{k=-\infty}^{\infty} \dot{I}_\alpha(t-kT; P_{0,T})\,. \end{aligned}$$

where we have used the same trick dealing with the derivative of $\mathcal{D}$ as is used to handle the derivative of the Delta function.

Similarly we can show, analogously to what we previously found for a $\theta(t)$ perturbation, that the frequency dependence of $\sigma_{\alpha y}$ can be extracted by

$$\sigma_{\alpha y}(\omega) = \lim_{\delta\to 0} \frac{T\, s}{1-e^{-sT}} \int_0^{\infty} \mathrm{d}t\, e^{-st}\, I_\alpha(t; P_{0,T})\,.$$