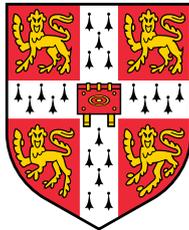

Nonzero temperatures and emergent disorder in spin liquids

OLIVER COLIN MALLINSON HART



*Corpus Christi College
University of Cambridge*

SUPERVISOR:
PROF. CLAUDIO CASTELNOVO

JUNE 2021

This dissertation is submitted for the degree of
Doctor of Philosophy

To my parents

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface, Acknowledgements, and as specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. This thesis does not exceed 60 000 words.

Oliver Hart

June 2021

Abstract

Nonzero temperatures and emergent disorder in spin liquids

OLIVER COLIN MALLINSON HART

At absolute zero temperature, spin liquids are known to exhibit a fascinating array of phenomena including topological order, emergent gauge fields, and the existence of exotic quasiparticles with fractional statistics. Meanwhile, at temperatures that are high compared to all characteristic interaction energy scales in the system, they often behave as trivial, uncorrelated paramagnets. This thesis aims to better understand the behaviour of spin liquids at temperatures that are intermediate, lying between these two extremes. At such intermediate temperatures, a finite density of defects are thermally excited, which inherit the peculiar properties of the proximate, zero temperature spin liquid state. Although the two limits are in most cases continuously connected via a crossover, these exotic defects have the potential to drastically alter the properties of the system. In spite of half a century of progress since the original ideas of Pauling and Anderson, quantum spin liquids have evaded unambiguous experimental detection. Therefore any precursor diagnostics in this temperature regime would be greatly beneficial before attempting to reach challengingly low temperatures where the effect of thermally excited defects is altogether negligible.

The thesis begins by considering the evolution of quantum mechanical entanglement with temperature in the toric code, utilising concepts borrowed from the field of quantum information theory. This highlights the importance of an intermediate temperature regime in which the ground state manifold does not have well-developed quantum coherence, but excitations out of the ground state sector are exponentially suppressed with temperature by their gap. The focus then moves to understanding how high-energy quasiparticles in the toric code propagate through a background of incoherent, thermally-generated gauge field excitations within this intermediate temperature regime. Analytical progress is made by making use of a mapping to the Bethe lattice, allowing predictions to be made about the experimentally measurable finite temperature dynamical structure factor. I then look at the propagation of quasiparticles in a different context: Classical spin ice. When spin ice is subjected to rapid cooling, the density of excitations (monopoles) form long-lived metastable

plateaux as a result of the formation of noncontractible pairs. I develop a thorough understanding of the origin of these plateaux by formulating the problem in terms of reaction–diffusion processes and performing large-scale simulations, which both suggest that the long-range nature of the interactions between monopoles is the linchpin of the plateaux. Finally, I consider the nonequilibrium dynamics of quantum spin liquids in closed systems at finite energy density. In this context, the emergent (self-generated) nature of the disorder plays a crucial role, allowing the system to generically exist in a superposition of different disorder realisations. I show that this property gives rise to power law decay of the dynamical structure factor and unbounded logarithmic growth of entanglement in time, just as for many-body localised systems, despite the existence of a mapping to free fermions.

Publications comprising this thesis

1. Oliver Hart and Claudio Castelnovo, *Entanglement negativity and sudden death in the toric code at finite temperature*, *Phys. Rev. B* **97**, 144410 (2018).
Some preliminary results pertaining to this paper, presented in Chapter 2, were previously submitted as a Part III project at the University of Cambridge. Specifically, this applies to Equations (2.3.16), (2.4.1) and their associated derivations only, which are presently in this thesis for contextualisation of the new results.
2. Oliver Hart, Yuan Wan and Claudio Castelnovo, *Coherent propagation of quasiparticles in topological spin liquids at finite temperature*, *Phys. Rev. B* **101**, 064428 (2020).
3. Oliver Hart, Marianne Haroche and Claudio Castelnovo, *Long-range Coulomb interactions and nonhydrodynamic behavior in thermal quenches in spin ice*, *Phys. Rev. B* **100**, 184411 (2019)
4. Oliver Hart, Sarang Gopalakrishnan and Claudio Castelnovo, *Logarithmic entanglement growth from disorder-free localisation in the two-leg compass ladder*, *Phys. Rev. Lett.* **126**, 227202, (2021)

Other works not included in this thesis

1. Oliver Hart, Garry Goldstein, Claudio Chamon and Claudio Castelnovo, *Steady-state superconductivity in electronic materials with repulsive interactions*, *Phys. Rev. B* **100**, 060508(R) (2019).
2. Oliver Hart, Yuan Wan and Claudio Castelnovo, *Correlation holes and slow dynamics induced by fractional statistics in gapped quantum spin liquids*, *Nat Commun* **12**, 1459 (2021).
3. Oliver Hart, Giuseppe De Tomasi and Claudio Castelnovo, *From compact localized states to many-body scars in the random quantum comb*, *Phys. Rev. Research* **2**, 043267 (Editors' Suggestion).

Acknowledgements

First and foremost, I would like to thank Claudio for his continued support and academic guidance throughout the duration of my PhD. He has always asked the right questions, whilst giving me the liberty to pursue my own ideas where appropriate. It has been a pleasure to work with him, and to experience his inspirational enthusiasm for physics.

I would also like to thank everyone with whom I have had the privilege of collaborating, in particular Claudio Chamon, Garry Goldstein, Sarang Gopalakrishnan, Giuseppe De Tomasi, and Yuan Wan.

I am indebted to the numerous members of TCM, both past and present, that have made the last three and a half years productive and enjoyable, both inside and outside office hours. Especially to my office mates, Max and Alice, for the countless hours of conversation and gallons of coffee consumed, and my previous office mates Attila, Petr, and Tycho. Thank you to Angela, Elis, Michael, Ivona, Stephen, Chris, Beñat, and many others, too numerous to list here, for making TCM such a fun and welcoming environment. A special thanks goes to Attila, Elis, and Max who proofread parts of this thesis at short notice.

I would also like to acknowledge those outside of TCM that have contributed towards making my time in Cambridge so enjoyable. Thank you to Glenn, Álvaro, and everyone else that I have had the pleasure of meeting at the various conferences and summer schools I have attended throughout my PhD. Also, to my housemates Tomi and Mary, who have made the various lockdowns throughout 2020–2021 not just bearable but enjoyable; their unwavering encouragement and friendship has been invaluable in making life outside of work so much fun. Above all, thank you to Tasha for standing by my side through everything that the world has thrown at us during the last couple of years.

Finally, I cannot thank enough my family, Julie, Jeremy and Will, for the love and support they have provided me with, not only during the PhD, but throughout my undergraduate years at Cambridge, and beyond. I would not be where I am today were it not for them letting me pursue my interests.

Contents

1	Introduction	15
1.1	Thesis outline	17
1.2	Quantum mechanical entanglement	20
1.2.1	Entanglement and the Schmidt decomposition	21
1.2.2	Entanglement in condensed matter systems	22
1.2.3	Mixed state entanglement	23
1.3	Temperature in quantum mechanical systems	24
1.3.1	Relaxation and thermalisation in open quantum systems	24
1.3.2	Thermalisation of closed quantum systems	26
1.4	A brief introduction to spin liquids	30
1.4.1	Kitaev’s toric code model	30
1.4.2	Ising \mathbb{Z}_2 lattice gauge theory	35
1.4.3	Classical spin ice	37
1.4.4	Quantum spin ice and U(1) lattice gauge theory	43
1.4.5	Kitaev’s honeycomb model	45
1.5	Localisation phenomena	50
1.5.1	Anderson localisation	50
1.5.2	Many-body localisation	58
1.5.3	Disorder-free localisation	61
2	Entanglement negativity in the toric code	63
2.1	Motivation	63
2.2	Mixed state separability	64
2.2.1	PPT criterion	64
2.2.2	The entanglement negativity	65
2.2.3	The replica method	66
2.3	Calculations	66
2.3.1	Star plaquette pair	68
2.3.2	Extended boundary	71

2.4	Discussion	79
3	Dynamics of spinons at finite temperature	83
3.1	Motivation	83
3.2	Hopping Hamiltonian and perturbation theory	84
3.2.1	Perturbation theory and effective Hamiltonians	86
3.2.2	Spinon vacuum	87
3.2.3	Effective Hamiltonian for an isolated spinon	88
3.2.4	Finite temperature	90
3.3	Lattice walks	92
3.3.1	Single-particle Green's function	92
3.3.2	Density evolution	94
3.3.3	Interpretation	95
3.4	Self-retracing paths: Generating functions	95
3.4.1	Closed walks	97
3.4.2	Open walks	98
3.4.3	Constrained closed walks	99
3.4.4	Nonreversing walks	100
3.5	Analytical results	103
3.5.1	Single spinon density of states	103
3.5.2	Single spinon Green's function	104
3.5.3	Spinon density profile	105
3.6	Numerical results	109
3.6.1	Finite temperature	111
3.7	Conclusions	112
4	Thermal quenches in classical spin ice	115
4.1	Motivation	115
4.2	Monte Carlo methods	116
4.2.1	Single spin flip dynamics	119
4.2.2	Waiting time Monte Carlo	119
4.3	Background and summary of results	121
4.4	Models	124
4.4.1	Classical spin ice	124
4.4.2	Charges on diamond lattice	126
4.5	Monte Carlo Simulations	127
4.5.1	Classical spin ice	127

4.5.2	Charges on diamond lattice	133
4.6	Summary and mean field modelling	135
4.6.1	Short-time dynamics	137
4.6.2	Truncated interactions	138
4.6.3	Long-range Coulomb interactions	141
4.7	Conclusions	145
5	Logarithmic growth of entanglement in a two-leg ladder	149
5.1	Motivation	149
5.2	Models and mappings	151
5.2.1	Exact solution	152
5.2.2	Anderson localisation	153
5.3	Entanglement growth	155
5.4	Dynamical structure factor	160
5.5	Exponentially weak dephasing	162
5.6	Discussion	164
6	Conclusions and outlook	167
6.1	Summary	167
6.2	Outlook and future directions	169
A	Diffusion on other lattices	173
A.1	Toric code on the kagome lattice	173
A.2	Triangular lattice	173
A.3	Honeycomb lattice	175
B	Higher order moments	179
C	Corrections to finite size scaling exponent	181
D	Double charge contribution	183
E	Compass model dualities	187
F	Bogoliubov solution to TFIM	191
G	Dephasing perturbation theory	193

1 | Introduction

The development of quantum mechanics in the early years of the 20th century created a paradigm shift in the way that we describe the world around us. In a system described by classical physics, it is possible, at least in principle, to specify simultaneously the positions and momenta of all constituent particles with arbitrarily high precision. The quantum mechanical description of the same system is profoundly different; the positions and momenta fundamentally cannot be specified with absolute precision, and must, instead, be described probabilistically. The system as a whole is characterised by its *wave function*, which encodes the probability distributions for the outcomes of all measurements that could be performed on the system. In addition to its important philosophical implications, the description of a physical system in terms of a probabilistic wave function gives rise to a plethora of intriguing and often counter-intuitive phenomena.

The evolution of the system's wave function in time is governed by Schrödinger's equation, which formally resembles a wave equation. This means that a particle described within quantum mechanics exhibits phenomena that were previously associated with classical waves, such as water or sound waves, the most prominent of which is arguably interference. For instance, Young's famous double slit experiment, in which waves are permitted to pass through two thin slits separated by a distance comparable to the wavelength, leads to a similar interference pattern whether the experiment is performed using water waves or single electrons [1] (at vastly different length and time scales). In the context of condensed matter, interference is responsible for the celebrated Aharonov–Bohm effect, and Anderson localisation, both of which will feature prominently in this thesis.

Another surprising prediction of quantum mechanics is the notion of entanglement, which has no classical analogue. As Schrödinger put it in his 1935 paper discussing entanglement, “I would not call that *one* but rather *the* characteristic trait of quantum mechanics” [2]. Quantum entanglement corresponds to a unique type of interdependency amongst a system's constituents, whereby a subsystem and its complement cannot be described independently from one another. Consequently, it is possible for the state of one subsystem to depend upon the measurement outcome of its complementary subsystem. As quantified by Bell's theorem [3], this special type

of correlation is stronger than any correlations that could be produced classically. In the field of quantum information, entanglement is now considered as a resource (“as real as energy” [4]) with utility in performing nonclassical tasks such as quantum cryptography and teleportation.

The above phenomena are contingent on the system being described by a pure state or wave function, which is appropriate when the system is isolated from any environmental influence. However, this is only ever an approximation, and some coupling to the environment is assuredly inevitable. When the environment is included quantum mechanically, the combined system plus environment should be described by a pure state. Even if the system and its environment are initially prepared in a separable (i.e., unentangled) state, entanglement between them will be generated irreversibly, establishing an arrow of time, however weak the coupling. The system and its surroundings are therefore generically entangled with one another. If the system and its surroundings have mutually equilibrated, then the reduced state of the system alone will be described by an impure, or mixed, quantum state (density matrix) corresponding to one of the standard ensembles from statistical mechanics, with the system inheriting the temperature of the environment*. As the temperature of the system’s surroundings is increased, the so-called purity of the density matrix is reduced, and the system becomes increasingly classical in nature as its various degrees of freedom become progressively incoherent.

In this thesis, we are primarily interested in the impact of nonzero temperatures on spin liquids – both classical and quantum – and related models. Spin liquids are an example of systems that do not develop conventional order as they are cooled down, arising naturally in the context of frustrated magnetism. Typically, the fundamental constituents in such systems are spinful degrees of freedom that interact antiferromagnetically on lattices that frustrate these interactions, i.e., the way that the spins are arranged in space prevents all local energies from being satisfied simultaneously, resulting in a macroscopically large number of equally favourable states at low energies. Canonical examples of such frustrated lattices are the triangular and kagome lattices in two dimensions, or the pyrochlore lattice in three dimensions. The large number of low-energy states means that, at low temperatures, the system is neither completely disordered like a gas, nor does it possess long-range order like a solid; instead, a spin liquid is interposed between the two, with partially developed order at low temperatures. As a result, classical spin liquids exhibit a host of intriguing collective phenomena, such as a macroscopic residual ground state entropy at low temperatures, emergent gauge fields, and fractionalised excitations that in some cases interact via an emergent long-ranged Coulomb interaction.

A quantum spin liquid may be obtained from a classical spin liquid by adding extra ‘transverse’ (noncommuting) terms to the system’s Hamiltonian. In this way, the transverse terms select particular superpositions of the classical spin configurations from the macroscopically degenerate space

*This assumes that the system is connected to a heat bath containing many more degrees of freedom than the system alone, such that the effect of the system on the state of the bath is negligible.

of ground states. These superpositions can have vastly different properties from the underlying classical spin configurations, including emergent lattice gauge theory descriptions and “topological order”, wherein the ground state(s) exhibits long-ranged entanglement and the system’s fractionised quasiparticles can obey statistics intermediate between bosonic and fermionic, amongst other phenomena.

The motivation for the introduction of nonzero temperatures in these systems is often introduced as being entirely practical; as we have argued, the presence of nonzero temperature is generally unavoidable in experiment, and it is therefore pertinent to ask what features survive at finite temperature. However, as we will see, this viewpoint is perhaps too pessimistic. Due to energy exchange with the environment, nonzero temperature implies a finite density of defects will be excited within the system. In quantum spin liquids, although the spin liquid “phase” is only strictly defined at zero temperature (often they are continuously connected, via a crossover, to the infinite temperature paramagnet*), the thermally excited defects inherit exotic properties from the proximate spin liquid state. These exotic defects have the ability to drastically alter the properties of the system – thermodynamic, spectroscopic, transport or otherwise – in a qualitative and distinctive manner, that should be accessible in the relevant experiments. This falls under the general paradigm of treating defects and disorder (in this case induced by thermal fluctuations) as a potential probe of exotic behaviour, as opposed to being a hindrance. As part of the quest to obtain reliable experimental signatures of quantum spin liquid behaviour, any precursor diagnostics at intermediate temperatures would be greatly beneficial before attempting to reach challengingly low temperatures where the effects of defects can be altogether neglected. In this thesis, we explore some of the consequences of these thermally excited defects on a case-by-case basis in a variety of settings.

1.1 Thesis outline

In the remainder of this chapter, we will introduce the core concepts that underpin many of the main results presented in this thesis. Specifically, we begin by introducing entanglement in the setting of condensed matter systems, and its generalisation to mixed quantum states. We then explain how an *open* quantum system reaches thermal equilibrium when connected to a large heat bath, with a focus on the relevant time scales. In this context, we also introduce the eigenstate thermalisation hypothesis (ETH) in order to explain how the predictions of statistical mechanics emerge in a closed, strongly interacting quantum system. In Sec. 1.4, we introduce some of the prototypical models of spin liquids, both classical and quantum, using them to outline the central phenomenology. Finally, we introduce localisation phenomena in Sec. 1.5, which play an important role in finite temperature states possessing a finite density of defects, beginning with Anderson localisation in free fermion

*One notable exception to this statement is the toric code in three spatial dimensions [5], which exhibits a continuous, topological phase transition at $T > 0$ [6].

systems (and non-interacting spin chains that map to free fermions). Its generalisation to interacting systems, many-body localisation, which provides a robust mechanism for violating the ETH, is then discussed in Sec. 1.5.2. We close this chapter with a brief overview of how localisation phenomena can arise, at least over some appreciable time scale, in systems possessing translational invariance through the mechanism of disorder-free localisation.

In Chapter 2, we study the evolution of entanglement with temperature in a particular (exactly solvable) \mathbb{Z}_2 spin liquid, the toric code. We present the first exact calculation of the entanglement negativity – a measure of mixed state entanglement that is notoriously difficult to calculate in a many-body setting – at finite temperature in a two-dimensional lattice system. We show that the entanglement content of the system exhibits the phenomenon of ‘sudden death’: above a critical temperature the entanglement vanishes discontinuously, signalling that the system’s correlations can be entirely reproduced by a classical system at and above this temperature. This sudden death temperature depends on the length of the boundary between the two subsystems under consideration. These features are summarised by constructing an entanglement phase diagram for the system.

In Chapter 3, we focus on models where quantum spin liquid behaviour is facilitated by perturbative ring-exchange processes. In such models there exists a natural separation of energy scales between the different species of quasiparticle excitation. This permits an intermediate temperature regime in which the ground state manifold does not have well-developed quantum coherence, but excitations *out* of the ground state sector are exponentially suppressed with temperature by their gap. Via the mechanism of disorder-free localisation, the system generates its own disorder and maps onto a tight-binding model subject to a fluctuating emergent \mathbb{Z}_2 gauge field; a direct consequence of the anyonic statistics of the excitations. Using the approximation of self-retracing paths, we are able to study analytically how the high-energy quasiparticles propagate through the background of incoherent gauge field excitations, which manifests in experimentally-measurable quantities such as the finite-temperature dynamical structure factor.

In Chapter 4, we shift gear and look instead at some of the intriguing phenomena that can happen at finite temperature in *classical* spin liquids. Specifically, we look at thermal quenches in the context of classical spin ice. Classical spin ice hosts excitations known as magnetic monopoles that are born out of the underlying spins. These excitations behave in many respects as emergent particles that interact via an effective long-ranged Coulomb interaction. The monopoles do however differ from free particles in that they are able to form *noncontractible pairs* – a pair of oppositely-charged monopoles separated by a spin whose reversal does not lead to their annihilation. Such pairs of monopoles are metastable, requiring energy to separate the constituent quasiparticles and annihilate them elsewhere on the lattice. At low temperatures, and in the absence of stray monopoles, these noncontractible pairs are long-lived and hence it is favourable to treat them as a distinct species of quasiparticle. Phrasing the problem in terms of reaction-diffusion processes, we are able to understand more quantitatively the origin of the metastable plateaux in monopole density that

develop after cooling classical spin ice rapidly. We find that the long-range nature of the interactions between monopoles is fundamental in driving the formation of the plateaux, which is confirmed by employing large-scale Monte Carlo simulations.

Chapter 5 is again concerned with the behaviour of systems where quantum coherence plays a significant role. We look in particular at a one-dimensional spin ladder known as the compass model, relevant to the description of transition metal oxides, which hosts an extensive number of local, \mathbb{Z}_2 valued conserved quantities. As in Chapter 3, these conservation laws allow the system to generate its own quenched disorder through the mechanism of disorder-free localisation. Consequently, the system is Anderson localised in a typical symmetry sector at finite energy density above the ground state, and behaves in many respects as though it were subjected to quenched disorder. However, certain quantities are acutely sensitive to the emergent character of the disorder, and will be our primary focus. We show that starting from generic initial states, entanglement grows logarithmically; in addition, equilibrium dynamical correlation functions decay with an exponent that varies continuously with temperature and model parameters. These features are commonly attributed to the phenomenon of many-body localisation. We elucidate how such features arise from a common origin – the self-generated nature of the disorder – and how *single particle* localisation can give rise to the same phenomenology as *many-body* localisation.

1.2 Quantum mechanical entanglement

In a *classical* closed many-body system of interacting particles, the positions and momenta of all particles can be specified simultaneously with arbitrary precision. This implies that if the system is partitioned into multiple, disjoint subsystems, the positions and momenta in one subsystem can be specified independently from the other complementary subsystems. The same is not necessarily true in isolated quantum mechanical systems, which permit a new type of interdependency between the constituent degrees of freedom known as quantum mechanical *entanglement*.

An isolated quantum mechanical system is described by a pure state $|\Psi\rangle$, which specifies the probability distribution of outcomes of all measurements that can in principle be made on the system. When the system is partitioned into two disjoint subsystems, the probability distributions pertaining to the two individual subsystems are not necessarily independent from one another in the statistical sense. The system is quantum mechanically entangled if the two probability distributions are interconnected, i.e., if the probability distribution of measurement outcomes for the second subsystem depends on the outcome of the measurement performed on the first subsystem.

This feature is most crisply illustrated in a minimal system composed of two spin-1/2 degrees of freedom, each considered as its own subsystem. One famous example of a maximally entangled state of the two spins is one of the four Bell states

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle_1 \otimes |\uparrow\rangle_2 + |\downarrow\rangle_1 \otimes |\downarrow\rangle_2) , \quad (1.2.1)$$

where $|\uparrow\rangle$ (“spin up”) and $|\downarrow\rangle$ (“spin down”) are eigenstates of $\hat{\sigma}^z$ with eigenvalues $+1$ and -1 , respectively. If spin one is measured first, one will observe the measurement outcomes “up” and “down” with equal probability. If spin two is measured immediately after the first, the outcome of the second measurement will always be equal to the outcome of the first measurement (i.e., with unit probability). This behaviour is to be contrasted with the separable state

$$|\Psi\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle_1 + |\downarrow\rangle_1) \otimes (|\uparrow\rangle_2 + |\downarrow\rangle_2) , \quad (1.2.2)$$

which possesses no quantum mechanical entanglement. Now, if spin one is measured first, one again observes the measurement outcomes “up” and “down” with equal probability. However, the measurement of the second spin will *also* lead to measurement outcomes “up” and “down” with equal probability, *irrespective* of the outcome of the first measurement.

1.2.1 Entanglement and the Schmidt decomposition

More generally, in a system $S = A_1 \cup A_2$, a state $|\psi\rangle \in \mathcal{H}_{A_1} \otimes \mathcal{H}_{A_2}$ is known as *separable* (i.e., not entangled) if it can be written as a product state over the tensor product Hilbert space

$$|\psi\rangle = |\phi\rangle_1 \otimes |\chi\rangle_2, \quad (1.2.3)$$

where $|\phi\rangle_1 \in \mathcal{H}_{A_1}$ and $|\chi\rangle_2 \in \mathcal{H}_{A_2}$ [7]. Physically, this property of $|\psi\rangle$ means that any measurements made on the composite system $A_1 \cup A_2$ separates into the product of the expectation values on the two subsystems individually: $\langle \hat{A} \otimes \hat{B} \rangle = \langle \hat{A} \otimes \mathbb{1} \rangle \langle \mathbb{1} \otimes \hat{B} \rangle$, and measurements on one subsystem do not affect the quantum state that describes the other subsystem.

An *entangled* state is one that cannot be written in the form (1.2.3), i.e., one cannot attribute a single quantum mechanical state to each of the two subsystems [7]. For entangled states, the outcomes of measurements made on the composite system are correlated so that $\langle \hat{A} \otimes \hat{B} \rangle \neq \langle \hat{A} \otimes \mathbb{1} \rangle \langle \mathbb{1} \otimes \hat{B} \rangle$. The entanglement properties of a given state $|\psi\rangle$ can be examined using its Schmidt (singular value) decomposition [8] into states $\{|\psi_i\rangle_1\}$ and $\{|\psi_i\rangle_2\}$ on subsystems A_1 and A_2 , respectively

$$|\psi\rangle = \sum_i \sqrt{p_i} |\psi_i\rangle_1 \otimes |\psi_i\rangle_2. \quad (1.2.4)$$

The real, positive coefficients $\sqrt{p_i}$ are known as the Schmidt coefficients. Normalisation of the state $|\psi\rangle$ implies that $\sum_i p_i = 1$. The decomposition is unique up to “rotations” within any degenerate subspaces that the Schmidt spectrum may contain [7]. Separability corresponds to the case that exactly one of the Schmidt coefficients is nonzero, in which case we reobtain (1.2.3). Conversely, for a maximally entangled state, all Schmidt coefficients are equal.

The reduced density matrices of subsystems A_1 and A_2 are found by tracing out the degrees of freedom belonging to the complementary subsystem. One may show from (1.2.4) that

$$\hat{\rho}_{A_1} = \text{Tr}_{A_2} \hat{\rho} = \sum_i p_i |\psi_i\rangle_1 \langle \psi_i|_1, \quad (1.2.5)$$

$$\hat{\rho}_{A_2} = \text{Tr}_{A_1} \hat{\rho} = \sum_i p_i |\psi_i\rangle_2 \langle \psi_i|_2. \quad (1.2.6)$$

The two reduced density matrices have the same spectrum $\{p_i\}$. When the state $|\psi\rangle$ is separable, the reduced density matrices $\hat{\rho}_{A_1}$ and $\hat{\rho}_{A_2}$ are pure, $\text{Tr} \hat{\rho}_{A_i}^2 = 1$. A natural measure of the separability of $|\psi\rangle$ is the von Neumann entropy of the reduced state $\hat{\rho}_{A_1}$ [8, 9]

$$S(\psi) = -\text{Tr} \{\hat{\rho}_{A_1} \ln \hat{\rho}_{A_1}\} = -\sum_i p_i \ln p_i. \quad (1.2.7)$$

A pure state has $S(\psi) = 0$, while a maximally entangled state has $S(\psi) = \ln D$, where $D = \min(D_1, D_2)$ and $D_i = \dim \mathcal{H}_{A_i}$. Supposing that both A_1 and A_2 both contain N spin-1/2 degrees of freedom, a maximally entangled state has $S(\psi) = N \ln 2$.

1.2.2 Entanglement in condensed matter systems

Up to this point, the discussion has been completely general, and applies to any pure quantum state. However, in the context of condensed matter, we are often interested in the low temperature properties of a given physical system. We are therefore seldom concerned with the behaviour of arbitrary pure states; rather, the ground state(s) of the system's Hamiltonian typically take centre stage. Such Hamiltonians are often *local*, and this locality is reflected in the behaviour of the entanglement entropy content of the ground state. Focusing on one spatial dimension, for systems with a finite correlation length (i.e., excluding, for now, critical systems) the ground state obeys the so-called *area law*

$$S(\ell) \propto \ell^{d-1} = \text{const.}, \quad (1.2.8)$$

independent of ℓ^* , the volume (i.e., length for $d = 1$) of the subregion under consideration. This result was first observed by Calabrese and Cardy in Ref. [10], and later proven by Hastings [11]. In 1D critical systems, the area law is replaced by a logarithmic increase of the entanglement entropy with subsystem size

$$S(\ell) = \frac{c}{3} \log \frac{\ell}{a} + \dots, \quad (1.2.9)$$

where c is the central charge, which characterises the universality class of the conformal field theory that describes the critical point, and a is an ultraviolet (i.e., lattice scale) cutoff. The dots corresponds to nonuniversal corrections. In higher dimensions, $d > 1$, the situation is more intricate, and we have fewer exact techniques at our disposal to study such systems rigorously. For the ground states of *gapped* systems comprised of free fermions or free bosons, one can show that the entanglement entropy satisfies an area law [12], whilst for free fermions possessing a well-defined Fermi surface, the area law is weakly violated by means of a logarithmic enhancement, $S(\ell) \sim \ell^{d-1} \ln \ell$ [13].

The entanglement entropy is also of fundamental *practical* importance for characterising the feasibility of simulating one dimensional systems numerically. This stems from a link between low-entanglement, area law ground states and so-called matrix product states (MPS) [14–16], wherein the wave function coefficients are compactly and efficiently parametrised in terms of a product of finite-dimensional matrices. It is possible to show that a one dimensional quantum state can be efficiently represented by an MPS if it satisfies an area law [12]. This result underlies the success of the density matrix renormalisation group (DMRG) [17], and other tensor network based

*This result is only independent of subsystem size as long as the size of both subsystems significantly exceeds the correlation length. The value of the constant is set approximately by the correlation length ξ , with $S \sim \log \xi$.

techniques, in one spatial dimension.

1.2.3 Mixed state entanglement

At finite temperature (as we will discuss later), or when considering partitions of the system into more than two disjoint components, it becomes necessary to describe the system using a mixed state $\hat{\rho}$. For example, in the case of a *tripartition* of the system $S = A_1 \cup A_2 \cup A_3$, the subsystems A_1 and A_2 are described by the (in general) mixed state $\hat{\rho}_{12} = \text{Tr}_3 |\psi\rangle\langle\psi|$, if the system as a whole is described by the pure state $|\psi\rangle$. Separability over the two subsystems A_1 and A_2 is then defined in the following way: a mixed state $\hat{\rho}$ is separable if it can be written in the form [8]

$$\hat{\rho} = \sum_i p_i \hat{\rho}_i^{(1)} \otimes \hat{\rho}_i^{(2)}. \quad (1.2.10)$$

The weights p_i are real, positive and normalised $\sum_i p_i = 1$, i.e., they represent probabilities. The two subsystems A_1 and A_2 are no longer uncorrelated, but the correlations are “classical” in nature, in the sense that expectation values of factorisable operators $\hat{A} \otimes \hat{B}$ assume the form

$$\langle \hat{A} \otimes \hat{B} \rangle = \sum_i p_i \langle \hat{A} \rangle_i \langle \hat{B} \rangle_i, \quad (1.2.11)$$

where $\langle \hat{A} \rangle_i = \text{Tr}[\hat{\rho}_i^{(1)} \hat{A}]$, and similarly for $\langle \hat{B} \rangle_i$. Such states satisfy Bell’s inequality [18, 19], and can easily be created using only local operations and classical communication [8]. Analogously to the case of pure states, an *entangled* mixed state is defined in the negative sense as any state that cannot be written in the form (1.2.10).

1.3 Temperature in quantum mechanical systems

1.3.1 Relaxation and thermalisation in open quantum systems

In any realistic experimental setting, a quantum system is never truly isolated from its surroundings (which we will interchangeably refer to as its environment). This nonzero coupling between a system and its environment will generically lead to decoherence and the creation or decay of excitations in the system (dissipation). If the system and the environment are coupled only weakly, then the state of the system can remain approximately pure for a long time (we will quantify how long this time scale can be later). Eventually, however weak the coupling, the system and the environment will asymptotically become entangled and the state of the system will become mixed. For sufficiently weak system–environment couplings, the asymptotic, steady-state density matrix is diagonal in the eigenbasis of the system [20]. If the environment is in equilibrium at temperature $T = \beta^{-1}$ ($k_B = \hbar = 1$ throughout this thesis), and certain ergodicity requirements are satisfied [21–23], then the steady-state density matrix corresponds to the canonical ensemble from statistical mechanics, $\hat{\rho} \propto e^{-\beta \hat{H}_S}$, where \hat{H}_S is the system Hamiltonian*.

Suppose that we are interested in the properties of a system S described by the Hamiltonian \hat{H}_S , which is weakly coupled to a reservoir B with Hamiltonian \hat{H}_B . The Hamiltonian of the combined system is of the general form

$$\hat{H} = \hat{H}_S + \hat{H}_B + \hat{H}_I, \quad (1.3.1)$$

where \hat{H}_I describes the system–reservoir interaction. In the Schrödinger picture, the time evolution of the combined system’s density matrix $\hat{\rho}(t)$ is determined by the Liouville–von Neumann equation

$$\frac{d\hat{\rho}(t)}{dt} = -i[\hat{H}, \hat{\rho}(t)], \quad (1.3.2)$$

In general, we would like to know how the state of the system S evolves in time. This is quantified by its reduced density matrix $\hat{\rho}_S(t) = \text{Tr}_B \hat{\rho}(t)$, obtained by tracing out the environment from the full density matrix $\hat{\rho}(t)$. An initially uncorrelated state of the total system, $\hat{\rho}(0) = \hat{\rho}_S(0) \otimes \hat{\rho}_B(0)$, will then evolve in time according to the dynamical map \mathcal{V}_t [24]

$$\hat{\rho}_S(t) = \mathcal{V}_t \hat{\rho}_S(0) \equiv \text{Tr}_B \left\{ \hat{U}(t) [\hat{\rho}_S(0) \otimes \hat{\rho}_B(0)] \hat{U}^\dagger(t) \right\}, \quad (1.3.3)$$

where $\hat{U}(t)$ is the time evolution operator in the Schrödinger picture. Assuming weak coupling between the system and the environment of the form $\hat{H}_I = \sum_\alpha \hat{A}_\alpha \otimes \hat{B}_\alpha$, and further assuming that

*In general, in the presence of stronger system–environment interactions, the system’s steady state – if such a state exists – will depend nontrivially on the interaction and bath Hamiltonians. Hereafter, we will consider only the case in which the system and its environment are coupled weakly.

memory effects in the bath can be neglected (the Markov approximation), it is possible to obtain a microscopic expression for (1.3.3), leading to the so-called quantum master equation (see, e.g., Refs. [7, 23, 24], which inspired this discussion, for a complete derivation)

$$\frac{d\hat{\rho}_S(t)}{dt} = \sum_{\omega, \omega'} \sum_{\alpha, \beta} e^{i(\omega' - \omega)t} \Gamma_{\alpha\beta}(\omega) \left(\hat{A}_\alpha(\omega) \hat{\rho}_S(t) \hat{A}_\beta^\dagger(\omega') - \hat{A}_\alpha^\dagger(\omega') \hat{A}_\beta(\omega) \hat{\rho}_S(t) \right) + \text{H.c.} \quad (1.3.4)$$

where $\hat{\rho}_S(t) = e^{i(\hat{H}_S + \hat{H}_B)t} \hat{\rho}_S(t) e^{-i(\hat{H}_S + \hat{H}_B)t}$ is the system's reduced density matrix in the interaction picture. The coefficients $\Gamma_{\alpha\beta}(\omega)$ are related to the Fourier transform of the bath correlation functions

$$\Gamma_{\alpha\beta}(\omega) = \int_0^\infty d\tau e^{i\omega\tau} \langle \hat{B}_\alpha^\dagger(t) \hat{B}_\beta(t - \tau) \rangle_{\hat{\rho}_B}, \quad (1.3.5)$$

while the operators $\hat{A}_\alpha(\omega)$ are related to the Fourier decomposition of $\hat{A}_\alpha(t)$, where the time dependence is again in the interaction picture:

$$\hat{A}_\alpha(t) = e^{i\hat{H}_S t} \hat{A}_\alpha e^{-i\hat{H}_S t} = \sum_{\omega} e^{-i\omega t} \sum_{\epsilon_n - \epsilon_m = \omega} |m\rangle \langle m| \hat{A}_\alpha |n\rangle \langle n| \equiv \sum_{\omega} e^{-i\omega t} \hat{A}_\alpha(\omega). \quad (1.3.6)$$

The states $|m\rangle$ and $|n\rangle$ in the above expression are eigenstates of \hat{H}_S with eigenvalues ϵ_m and ϵ_n , respectively. If one additionally applies the rotating wave approximation^{*}, then equation (1.3.4) can be written as a dynamical map[†],

$$\frac{d\hat{\rho}_S(t)}{dt} = -i[\hat{H}_L, \hat{\rho}_S] + \sum_{\omega} \sum_{\alpha, \beta} \gamma_{\alpha\beta}(\omega) \left(\hat{A}_\beta(\omega) \hat{\rho}_S \hat{A}_\alpha^\dagger(\omega) - \frac{1}{2} \left\{ \hat{A}_\alpha^\dagger(\omega) \hat{A}_\beta(\omega), \hat{\rho}_S \right\} \right), \quad (1.3.7)$$

where the so-called Lamb-shift Hamiltonian

$$\hat{H}_L = \sum_{\omega} \sum_{\alpha, \beta} S_{\alpha\beta}(\omega) \hat{A}_\alpha^\dagger(\omega) \hat{A}_\beta(\omega), \quad (1.3.8)$$

leads to a renormalisation of the system energy levels due to interactions with the environment, and $\Gamma_{\alpha\beta}(\omega) = \frac{1}{2} \gamma_{\alpha\beta}(\omega) + i S_{\alpha\beta}(\omega)$, such that $(S_{\alpha\beta}) \gamma_{\alpha\beta}$ is (anti-) Hermitian. If the bath is in a Gibbs state $\hat{\rho}_B = e^{-\beta \hat{H}_B} / Z_B$, where the partition function Z_B ensures normalisation of $\hat{\rho}_B$, then it can be

^{*}This approximation is *not* generically satisfied in condensed matter systems. In order for the rotating wave approximation to provide an accurate description of the dynamics (over coarse-grained time scales) the system should have a discrete spectrum, and the relaxation time scale of the system S, denoted τ_r , should satisfy $\tau_r \gg |\omega - \omega'|^{-1}$ [23, 24]. One can however derive an analogous master equation in the so-called singular coupling limit [23], in which the rotating wave approximation is not necessary.

[†]A dynamical map must satisfy the conditions of complete positivity, convex linearity and trace preservation [24].

shown that for any initial state the system will relax towards the Gibbs ensemble

$$\hat{\rho}_S(t) \xrightarrow{t \rightarrow \infty} \frac{1}{Z_S} e^{-\beta \hat{H}_S}, \quad (1.3.9)$$

as long as certain ergodicity requirements* are satisfied [21–23]. The (positive) eigenvalues $\gamma_n(\omega)$ of the Hermitian matrix $\gamma_{\alpha\beta}(\omega)$ have dimensions of inverse time (if the operators \hat{B}_α are chosen to be dimensionless), and can be interpreted as decoherence rates, which dictate how quickly the system approaches its stationary state given by (1.3.9).

We have therefore shown that a quantum system weakly coupled to a large heat bath, under a number of assumptions, relaxes towards the Gibbs state (or canonical ensemble) $\hat{\rho}_S \propto e^{-\beta \hat{H}_S}$ in accordance with the principle of maximum entropy, where the system inherits the temperature T of the heat bath.

1.3.2 Thermalisation of closed quantum systems

In the above, we have shown that thermalisation occurs under rather general conditions if the system of interest is connected to an external reservoir. However, we have implicitly assumed that the heat bath is itself prepared in a (mixed) thermal Gibbs state $\hat{\rho}_B \propto e^{-\beta \hat{H}_B}$. In order to describe how the *bath* reaches thermal equilibrium, we are led to consider the apparent paradox of thermalisation in closed quantum systems, and the recent drive to understand it (see, e.g., Refs. [25–29] for reviews). That is, understanding the mechanism by which an isolated quantum system is able to reach an equilibrium state, determined by only a small number of parameters such as temperature, chemical potential, etc., irrespective of the (pure) state in which the system was prepared. This statement appears to be at odds with the unitarity of quantum mechanics, since the unitary time evolution of the wave function does not “remove” any of the information about the system’s initial conditions. In the following, we will explain the resolution of this apparent paradox with a powerful conjecture about the properties of individual eigenstates known as the eigenstate thermalisation hypothesis (ETH) [30–32]. Conversely, systems that do not obey the ETH retain memory of their initial conditions. Localisation, which we will introduce later in Sec. 1.5, appears to be the only *robust* way in which a system is able to evade thermalisation[†].

In order to discuss the thermalisation of a closed quantum system, it is necessary to introduce the concept of a quantum quench. The system is initially prepared in some state $|\psi(0)\rangle$, which is often

*Specifically, the operators \hat{A}_α must connect all of the energy levels of the system Hamiltonian. This condition is ensured if there exist no operators acting on S that simultaneously commute with all the \hat{A}_α and \hat{H}_S [21, 22]. If the converse is true and such an operator *does* exist, then the Hilbert space of the system can be partitioned into disconnected sectors.

[†]Integrable systems, which possess an extensive number of conserved quantities, and undergo relaxation towards a so-called generalised Gibbs ensemble [33–35], are excluded from this statement as their integrability is not preserved upon addition of *generic* perturbations to the Hamiltonian.

the ground state of some preparation Hamiltonian \hat{H}_0 . For times $t > 0$, the state $|\psi(0)\rangle$ is evolved under unitary dynamics with respect to a *different* Hamiltonian \hat{H} , i.e., $|\psi(t)\rangle = e^{-i\hat{H}t} |\psi(0)\rangle$. If the state $|\psi(0)\rangle$ is not an eigenstate of \hat{H} , then its dynamics will be nontrivial. Expanding the initial state in the eigenbasis $\{|\Psi_n\rangle\}$ of \hat{H} , $|\psi(0)\rangle = \sum_n \langle\Psi_n|\psi(0)\rangle |\Psi_n\rangle \equiv \sum_n c_n |\Psi_n\rangle$, the long-time average of any observable \hat{O} , defined by $\overline{O} \equiv \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau O(\tau)$, may be written as

$$\overline{\langle\hat{O}\rangle} = \sum_n |c_n|^2 O_{nn}, \quad (1.3.10)$$

where O_{nn} are the matrix elements of the operator \hat{O} in the eigenbasis of \hat{H} , and the spectrum is assumed to be free of any degeneracy. The off-diagonal matrix elements oscillate in time and, hence, are removed by the time average. The time-averaged expectation value (1.3.10) can alternatively be viewed as an average with respect to a density matrix known as the ‘‘diagonal ensemble’’, $\overline{\langle\hat{O}\rangle} = \text{Tr}[\hat{O}\hat{\rho}_{\text{diag}}]$, where $\hat{\rho}_{\text{diag}} = \sum_n |c_n|^2 |\Psi_n\rangle\langle\Psi_n|$ [33, 36, 37]. Evidently, the state $\hat{\rho}_{\text{diag}}$, which describes the time-average of a completely generic observable, depends explicitly on the system’s initial conditions, and is therefore incompatible with the notion that the system’s density matrix should relax towards the appropriate equilibrium ensemble from statistical mechanics. Instead, as was realised in 1929 by von Neumann [38], one should focus not on the equilibration of the state of the system as a whole, but on the behaviour of a suitable set of ‘‘physical’’ observables that are accessible in experiment. Equivalence between the diagonal ensemble and the prediction of statistical mechanics is then established if the (time-averaged) expectation values in the two cases coincide for this privileged set of physical observables*. Operators such as the projector $|\Psi_n\rangle\langle\Psi_n|$, whose expectation value never relaxes, are disregarded on the basis that such operators are typically non-local, and, hence, eminently unphysical†.

In the 1990s, Srednicki introduced a generalisation of the random matrix theory prediction for the matrix elements of local operators in a series of pioneering papers [30–32]. The ansatz is known as the eigenstate thermalisation hypothesis, which takes the form [32]

$$O_{mn} = \mathcal{O}(E)\delta_{mn} + e^{-S(E)/2} f_O(E, \omega) R_{mn}, \quad (1.3.11)$$

where $E = (E_m + E_n)/2$ is the mean energy of the two eigenstates m and n , and $\omega = E_m - E_n$ is their energy difference. The first term, which describes the behaviour of the diagonal matrix elements, consists of a smooth function $\mathcal{O}(E)$ of the average energy. The second piece determines the decay of off-diagonal matrix elements. Specifically, $S(E) = \sum_n E\delta(E - E_n)$ is the thermodynamic entropy, $f_O(E, \omega)$ is a smooth function of E and ω , and the R_{mn} are random

*One also requires that the temporal fluctuations of the observable should vanish in the thermodynamic limit.

†Many-body localised systems, in which the l-bits (emergent conserved quantities that may be constructed from projectors onto eigenstates [39]) are local, are an exception to this intuition.

numbers with vanishing mean and unit variance. The characteristic scale of the entropy is set by the mean level spacing δ at energy E , i.e., $\delta \sim E e^{-S}$. Since the level spacing is typically exponentially small in the volume of the system, this implies, in turn, that the off-diagonal elements of \hat{O} are exponentially small in the system's volume with respect to the diagonal elements.

Using the ETH ansatz (1.3.11), we can now calculate the time-averaged expectation value of any “physical” observable

$$\overline{\langle \hat{O} \rangle} = \sum_n |c_n|^2 O_{nn} \simeq \mathcal{O}(E), \quad (1.3.12)$$

where we have neglected terms proportional to the second derivative of $\mathcal{O}(E)$ that arise due to finite energy fluctuations in the initial state (which are often subextensive in volume for local Hamiltonians [28, 36]). Remarkably, the smooth nature of the function $\mathcal{O}(E)$ ensures that the result (1.3.12) is independent of the initial conditions, encoded by the c_n , as long as they are only nonzero over a sufficiently narrow energy window. This means that the time-averaged expectation value agrees with the prediction of statistical mechanics for sufficiently well-behaved initial states. Moreover, the exponential decay of the off-diagonal matrix elements with volume means that equilibration to the value $\mathcal{O}(E)$ can occur on time scales much shorter than the inverse level spacing [28]. Equilibration to the value $\mathcal{O}(E)$ is not just established in the time-averaged sense; instead, $\langle \hat{O}(t) \rangle$ relaxes to its equilibrium value and then remains nearby thereafter (save for quantum revivals in systems of finite size). In fact, the temporal fluctuations of $\langle \hat{O}(t) \rangle$ about its average value can be shown to be exponentially suppressed in the volume of the system

$$\sigma_O^2 \equiv \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t d\tau [\langle \hat{O} \rangle^2] - (\overline{O})^2 = \sum_{m \neq n} |c_m|^2 |c_n|^2 |O_{mn}|^2 \leq \max_{m \neq n} |O_{mn}|^2 \sim e^{-S(E)}, \quad (1.3.13)$$

since the entropy is an extensive quantity. In contrast, the time-averaged *quantum* fluctuations of \hat{O} also decay with the size of the system, but algebraically in the volume [28].

The eigenstate thermalisation hypothesis addresses the apparent paradox of thermalisation in closed quantum systems in a particularly simple and elegant manner. Thermalisation under the dynamics of an ergodic Hamiltonian is encoded in the structure of the individual eigenstates, with each and every eigenstate being thermal. All of the information about the initial state of the system is never erased by the unitary time evolution, but (typically) becomes accessible only through highly nonlocal, unphysical operators. The ETH prediction for the matrix elements of physical observables has been verified numerically in a number of low-dimensional quantum systems including 2D hard-core bosons [36], interacting spin chains [40], the two dimensional transverse field Ising model [41], and many others (see, e.g., the references contained within Ref. [28]).

Recently, there has been a significant amount of interest in systems that (weakly) violate the ETH through the mechanism of quantum many body scarring (for a recent review see Ref. [42]).

In such systems, there exist a small number (typically measure zero) of mid-spectrum area law entangled states, which are decoupled from the rest of the spectrum. Since these atypical states are so few in number compared to the complementary thermal eigenstates, one might imagine that they are irrelevant to the equilibration properties of the models in question. However, in some cases such as the so-called PXP model [43], these atypical, area law entangled states have an anomalously large overlap with translationally invariant, experimentally accessible initial states. Indeed, the interest in weak ergodicity breaking was triggered by a surprising experiment in 2018 on a Rydberg-atom quantum simulator, which unveiled persistent, periodic revivals for some initial states [44]. In this way, systems possessing many body scars violate the ETH in the weak sense, since a *typical* infinite temperature state will still thermalise. During my PhD I have contributed to this field, showing that quantum many-body scarring can occur in certain random and periodic comb-like structures that host interacting fermionic degrees of freedom [45].

1.4 A brief introduction to spin liquids

In the following, we will discuss some of the defining features of spin liquids, both classical and quantum, by introducing some of the prototypical theoretical models, each of which has played a unique role in aiding our understanding of these fascinating phases of matter.

1.4.1 Kitaev’s toric code model

We will begin with arguably the simplest model of a gapped, \mathbb{Z}_2 quantum spin liquid: Kitaev’s toric code [46]. Despite appearing highly artificial at first sight, the toric code Hamiltonian appears as a natural limiting case of Kitaev’s honeycomb model [47] (introduced later in Sec. 1.4.5), which contains only nearest-neighbour interactions between spin degrees of freedom, and is approximately realised in certain iridates [48, 49] and other compounds with strong spin orbit coupling (see, e.g., Refs. [50, 51] for reviews). Recently, it was also realised that the toric code Hamiltonian emerges perturbatively through the mechanism of combinatorial gauge symmetry [52] in some nearest-neighbour spin models, allowing \mathbb{Z}_2 spin liquids to be embedded and studied on quantum annealers [53]. Notwithstanding its simplicity, the toric code exhibits many of the striking phenomena associated with quantum spin liquids in general, and therefore offers a natural setting in which to explore their consequences.

Hamiltonian

The toric code model is conventionally defined on a 2D square lattice composed of N sites with periodic boundary conditions (i.e., on a torus). Spin- $\frac{1}{2}$ degrees of freedom are located on each of the $2N$ bonds of the lattice and interact via the Hamiltonian

$$\hat{H}_{\text{TC}} = -J_e \sum_s \hat{A}_s - J_m \sum_p \hat{B}_p \equiv -J_e \hat{S} - J_m \hat{P}, \quad (1.4.1)$$

where the labels s and p denote the “stars” (sites) and “plaquettes” of the lattice (see Fig. 1.1), respectively, and $\hat{A}_s \equiv \prod_{i \in s} \hat{\sigma}_i^x$, $\hat{B}_p \equiv \prod_{i \in p} \hat{\sigma}_i^z$ [46]. $J_e, J_m > 0$ are the two coupling constants of the model. The subscripts e and m derive from a correspondence with \mathbb{Z}_2 lattice gauge theory, explored at the end of this section. As usual, $\hat{\sigma}_i = (\hat{\sigma}_i^x, \hat{\sigma}_i^y, \hat{\sigma}_i^z)$ denote the Pauli matrices, which describe the spin on bond i .

Ground states and topological degeneracy

All operators \hat{A}_s, \hat{B}_p in the Hamiltonian commute and so can be diagonalised simultaneously. The property $\hat{A}_s^2 = \mathbb{1} = \hat{B}_p^2$ implies that the operators \hat{A}_s and \hat{B}_p have eigenvalues ± 1 , and the states

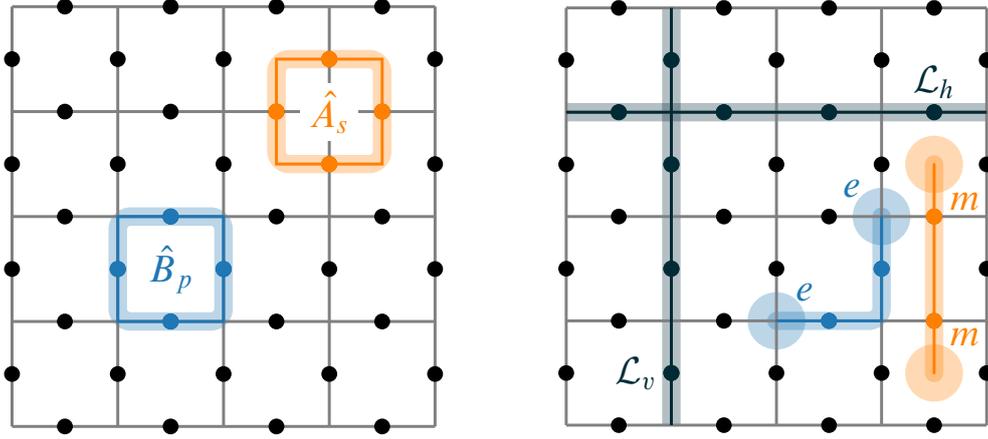


Figure 1.1: **Left:** Illustration of the “star” and “plaquette” operators, \hat{A}_s and \hat{B}_p , respectively, that appear in the toric code Hamiltonian (1.4.1). **Right:** Two possible noncontractible loops $\mathcal{L}_{h,v}$, and two open strings generating electric (e) and magnetic (m) defects at their ends. Recall that periodic boundary conditions are imposed. Electric (magnetic) defects inhabit the stars (plaquettes) of the square lattice, and cost an energy $2J_e$ ($2J_m$), so that each open string costs $4J_e$ ($4J_m$).

for which all \hat{A}_s , \hat{B}_p have eigenvalue $+1$ are the ground states. On a torus, the model has four ground states belonging to different “topological sectors,” which are classified according to the eigenvalues of non-local (system-spanning) operators. Such operators can be defined for instance as $\hat{\Gamma}_{h,v}^x = \prod_{i \in \mathcal{L}_{h,v}} \hat{\sigma}_i^x$ [54], where $\mathcal{L}_{h,v}$ are two noncontractible loops on the dual lattice that span the torus in the horizontal (h) and vertical (v) directions, respectively (see Fig. 1.1). It is not possible to distinguish between the four different topological sectors using any local operators. Similarly, no local operator can lead to a transition between the various ground states, which makes the toric code desirable from the perspective of encoding quantum information [46]. When the model is defined on a surface of genus g ($g = 1$ for the torus), the topological degeneracy is generalised to 4^g [55]. This dependence of the degeneracy on the *topology* of the underlying lattice is one of the many characteristic features of topologically ordered phases.

In order to write down an explicit expression for the ground state(s) of the model (1.4.1), it is convenient to introduce the group \mathcal{G} generated by products of star operators \hat{A}_s^* . The ground state of the model is then given by

$$|\psi_0\rangle = \frac{1}{|\mathcal{G}|^{1/2}} \sum_{g \in \mathcal{G}} g |0\rangle. \quad (1.4.2)$$

The state $|0\rangle$ is an eigenstate of $\otimes_i \hat{\sigma}_i^z$ belonging to the appropriate topological sector (i.e., it is also an eigenstate of the operators $\hat{\Gamma}_{h,v}^z$), and $|\mathcal{G}| = 2^{N-1}$. The state (1.4.2) is trivially an eigenstate of all \hat{B}_p since all group elements commute with \hat{B}_p , $[\hat{B}_p, g] = 0$, $\forall g \in \mathcal{G}$, and $\hat{B}_p |0\rangle = |0\rangle$. That

*The group elements $g \in \mathcal{G}$ must be defined modulo the identity for the inverse elements g^{-1} to be unique, since $\prod_s \hat{A}_s = \mathbb{1}$. That is, $g = \prod_s \hat{A}_s^{\sigma_s} = \prod_s \hat{A}_s^{1-\sigma_s}$ with $\sigma_s = 0, 1$. One must choose one of the two equivalent tilings of the \hat{A}_s operators for the mapping between g and $\{\sigma_s\}$ to be one-to-one.

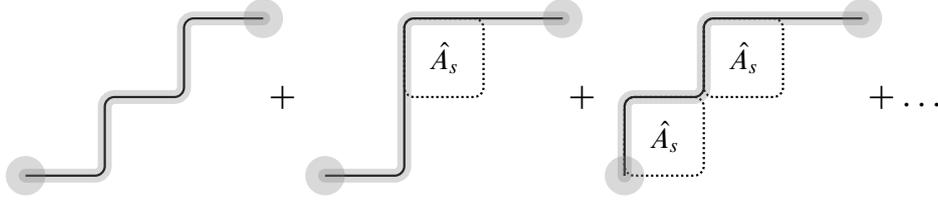


Figure 1.2: Schematic depiction of how every string ℓ' with the same endpoints is generated when acting with a given string ℓ of operators on the ground state.

(1.4.2) is an eigenstate of all \hat{A}_s follows from the fact that $\hat{A}_s \in \mathcal{G}$, giving $\hat{A}_s g = g'$, using the group properties of \mathcal{G} . Since the ground state corresponds to a sum over all group elements, we can relabel the mute index $g' \rightarrow g$ to give $\hat{A}_s |\psi_0\rangle = |\psi_0\rangle$, as required. Hence, the state (1.4.2) is an eigenstate of the Hamiltonian with energy $E_0 = -N(J_e + J_m)$.

Consider the expression for a given group element $g \in \mathcal{G}$, given as a product of star operators over some (possibly disconnected) region R : $g = \prod_{s \in R} \hat{A}_s$. Since all internal spins are flipped twice, the action of the group element g on the reference state $|0\rangle$ is to flip the spins on the *boundary* of R , denoted ∂R , so that we may write $g = \prod_{i \in \partial R} \hat{\sigma}_i^x$. Therefore, when expressed in the $\hat{\sigma}_i^z$ tensor product basis, as in (1.4.2), one observes that the ground state corresponds to an equal-weight superposition of closed loops of flipped spins. Note that it would be completely equivalent to formulate the ground state (1.4.2) in terms of the $\hat{\sigma}_i^x$ tensor product basis.

Fractionalisation and deconfinement of excitations

The toric code hosts two types of elementary excitations known as *electric charges* and *magnetic vortices*. These excitations are generated by *open* strings of Pauli matrices, $\hat{S}^\alpha(\ell) = \prod_{i \in \ell} \hat{\sigma}_i^\alpha$ with $\alpha = z, x$, respectively [55] (see Fig. 1.1). $\hat{S}^\alpha(\ell)$ flips the eigenvalues of the *two* stars (plaquettes) at the ends of the path ℓ , which lives on the direct (dual) lattice, costing an energy $4J_e$ ($4J_m$). When periodic boundary conditions are imposed, it is not possible to change the eigenvalue of just one star or plaquette operator as a result of the condition $\prod_s \hat{A}_s = \prod_p \hat{B}_p = \mathbb{1}$. Physically, flipping a spin changes the eigenvalue of the two adjacent stars (plaquettes)*.

Note that only the endpoints of the string matter; the choice of path ℓ is immaterial. Acting on the ground state (1.4.2) with, say, $\hat{S}^x(\ell)$, one obtains

$$\hat{S}^x(\ell) |\psi_0\rangle = \frac{1}{|\mathcal{G}|^{1/2}} \sum_{g \in \mathcal{G}} \hat{S}^x(\ell) g |0\rangle = \frac{1}{|\mathcal{G}|^{1/2}} \sum_{\ell'} \hat{S}^x(\ell') |0\rangle, \quad (1.4.3)$$

where the second summation is over all possible paths ℓ' with the same endpoints as the original

*This and the inability to change the eigenvalue of just one star/plaquette are in fact two separate statements. With open boundary conditions, it is not possible to locally create an unpaired excitation in the bulk, but the boundary can act as a source/sink of unpaired excitations.

path ℓ . This result is shown schematically in Fig. 1.2. Since the energy of a pair of defects is independent of their separation, the defects are *deconfined*, which is another characteristic feature of quantum spin liquids [51, 56].

At nonzero temperature $T = 1/\beta$, there is a finite probability of exciting defects given by the Boltzmann distribution $P(E) \propto e^{-\beta E}$. The exact density of species $X = e, m$ is

$$\rho_X = e^{-2\beta J_X} \frac{(1 + e^{-2\beta J_X})^{N-1} - (1 - e^{-2\beta J_X})^{N-1}}{(1 + e^{-2\beta J_X})^N + (1 - e^{-2\beta J_X})^N}. \quad (1.4.4)$$

In the limit of large N , the constraint that defects must be created in pairs becomes immaterial and the density is given approximately by the Fermi-Dirac distribution $n_F(2J_X) \equiv (e^{2\beta J_X} + 1)^{-1}$, for each species independently $X = e, m$.

Mutually semionic statistics

The excitations introduced in Sec. 1.4.1 obey mutually semionic statistics. This feature of topologically ordered phases is arguably the most important, and will play a prominent role in Chapter 3. As shown in Fig. 1.3, when one species encircles a quasiparticle of the *opposite* species, the wave function picks up a phase of π . Figure 1.3b differs from 1.3a by a closed loop of $\hat{\sigma}_i^x$ operators. When this loop contains an electric charge,

$$|\psi_b\rangle = \left(\prod_{i \in \partial R} \hat{\sigma}_i^x \right) |\psi_a\rangle = \left(\prod_{s \in R} \hat{A}_s \right) |\psi_a\rangle = -|\psi_a\rangle. \quad (1.4.5)$$

Since this operation can be interpreted as two successive particle exchanges, one may attribute a phase $\phi = \pi/2$ to exchanging electric charges and magnetic vortices [57]. In this way, the excitations are *mutually* semionic, but *individually* bosonic in the sense that interchanging two particles of the *same* species does not result in sign change of the wave function.

Entanglement entropy and stability

So far, we have seen that gapped \mathbb{Z}_2 quantum spin liquids (which by definition possess topological order) host deconfined excitations satisfying mutually semionic statistics, and admit a topological ground state degeneracy on a torus*. These interrelated features are also intimately connected to the entanglement properties of the ground state(s). As shown previously, the ground states of the toric code are an equal weight superposition of loops in either the σ_i^z or σ_i^x tensor product basis. The ground state entanglement, as measured by the von Neumann entanglement entropy S , takes the form of an area law with a subleading correction: $S \sim s_0 L - \gamma$, where s_0 is a constant, and L is the

*In fact the mutually semionic statistics of the two quasiparticle species, e and m , can be shown to *imply* a topological degeneracy [57, 58].

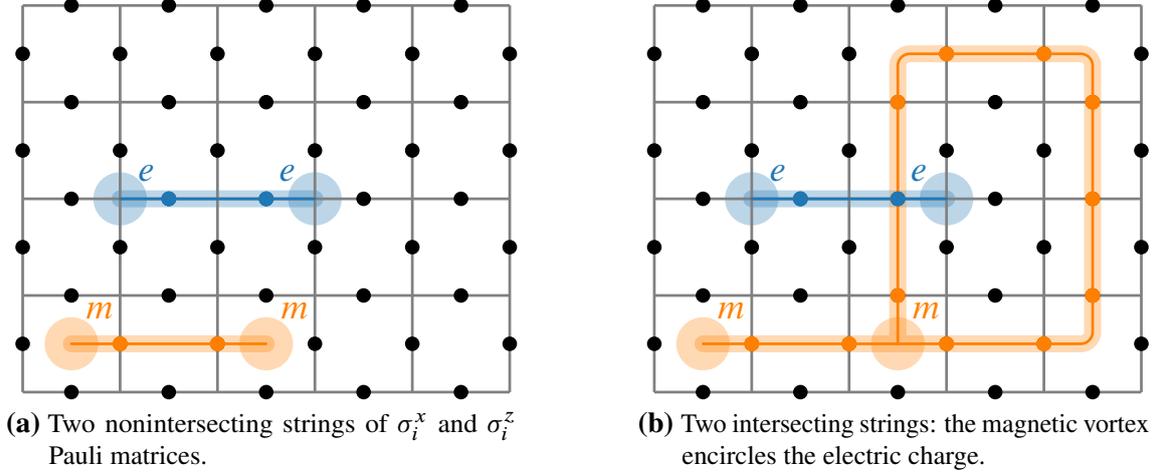


Figure 1.3: When a magnetic vortex encircles an electric charge, as in (b), the wavefunction picks up a phase of $e^{i\pi} = -1$ relative to configuration (a).

length of the (smooth [59, 60]) boundary separating two regions of the lattice (which are in direct contact). The universal correction term $\gamma = \ln 2$ is only nonzero in a system possessing topological order*, and thus can be used to identify topologically ordered phases of matter [51, 56].

The QSL phase of the toric code is robust to *all* sufficiently weak perturbations [63]. However, for sufficiently large perturbations, e.g., the application of a large magnetic field [64–67], the system undergoes a phase transition into a phase in which the excitations are confined. In the presence of a very large magnetic field h , say along z , the ground state is approximately the product state $|\sigma_i^z = 1\rangle$. Separating two charges in this state by flipping a string of spins costs an energy proportional to their separation and the strength of the field (each spin flip costs $\sim 2h$), implying confinement of the pair [56]. The confinement transition that occurs at some critical value of the applied field will be discussed in further detail in Sec. 1.4.2 in the context of Wegner’s Ising gauge theory.

Mapping to \mathbb{Z}_2 lattice gauge theory

The mapping to (Ising) lattice gauge theory (LGT) language can be formulated as follows [56]. The $\hat{\sigma}_i^z$ operator, which lives on the link ab , is written in terms of a vector potential $\hat{\sigma}_i^z = e^{i\hat{A}_{ab}}$, which has eigenvalues $A_{ab} = 0, \pi$. Similarly, the $\hat{\sigma}_i^x$ operators are written in terms of electric fields \hat{E}_{ab} , canonically conjugate to vector potential \hat{A}_{ab} , as $\hat{\sigma}_i^x = e^{i\pi\hat{E}_{ab}}$. In this way, $e^{i\hat{A}_{ab}}$ acts as a raising operator for the corresponding electric field, which assumes the values $E_{ab} = 0, 1$. The field variables are defined to be antisymmetric with respect to site indices, $\hat{A}_{ab} = -\hat{A}_{ba}$ and

*Note that there are other mechanisms that also give rise to subleading corrections to S , such as corner contributions. Indeed, an intricate subtraction scheme is in general necessary to isolate γ on a lattice [61, 62].

$\hat{E}_{ab} = -\hat{E}_{ba}$. With these definitions in place, the Hamiltonian (1.4.1) can be written in the form

$$\hat{H} = -J_e \sum_s e^{i\pi(\nabla \cdot \hat{E})_s} - J_m \sum_p \cos(\nabla \times \hat{A})_p. \quad (1.4.6)$$

The ‘‘lattice divergence,’’ $(\nabla \cdot \hat{E})_s$, and ‘‘lattice curl,’’ $(\nabla \times \hat{A})_p$, are given by the expressions

$$(\nabla \cdot \hat{E})_s \equiv \sum_{s' \in \langle ss' \rangle} \hat{E}_{ss'}, \quad \text{and} \quad (\nabla \times \hat{A})_p \equiv \sum_{b \in p} \hat{A}_{b, b+1}, \quad (1.4.7)$$

where the curl is evaluated in the clockwise direction. We identify $(\nabla \cdot \hat{E})_s$ with the charge on site s , and $(\nabla \times \hat{A})_p$ with the magnetic flux through the plaquette p , by analogy with Maxwell’s equations. The ground states are free of both electric charge and magnetic flux. The ground state sector for the electric charges is selected by the Gauss law constraint $(\nabla \cdot \hat{E})_s |\Psi\rangle = 0$. This mapping justifies the terminology ‘‘electric charges’’ and ‘‘magnetic vortices’’ used earlier in this section. The Hamiltonian is invariant under the unitary transformation

$$\hat{U}[\{\chi_s\}] = e^{i\pi \sum_s \chi_s (\nabla \cdot \hat{E})_s} = \prod_s \hat{A}_s^{\chi_s}, \quad (1.4.8)$$

generated by Gauss’ law, where $\chi_s = 0, 1$, under which the $\hat{\sigma}_i^z$ operators transform according to

$$\hat{\sigma}_{ss'}^z = e^{i\hat{A}_{ss'}} \rightarrow e^{i[\hat{A}_{ss'} + \pi(\chi_s - \chi_{s'})]}. \quad (1.4.9)$$

Evidently, the operators \hat{A}_{ab} transform like a gauge potential under a discrete gauge transformation, i.e., the analogue of $\mathbf{A} \rightarrow \mathbf{A} + \nabla \chi$ in classical electromagnetism. Note that the toric code exhibits a duality between its electric charges and magnetic vortices, and the above arguments could have equivalently been formulated with the roles of the electric field and vector potential swapped.

1.4.2 Ising \mathbb{Z}_2 lattice gauge theory

Here we consider a close relative of the toric code studied in the previous section: Wegner’s \mathbb{Z}_2 Ising lattice gauge theory [68], one of the simplest models that exhibits a phase transition between confining and topological (deconfined) phases. The Hamiltonian consists of the same term acting on the ‘‘stars’’ and a perturbing transverse magnetic field [68, 69]:

$$\hat{H}_{\mathbb{Z}_2} = -J \sum_s \hat{A}_s - h \sum_i \hat{\sigma}_i^z. \quad (1.4.10)$$

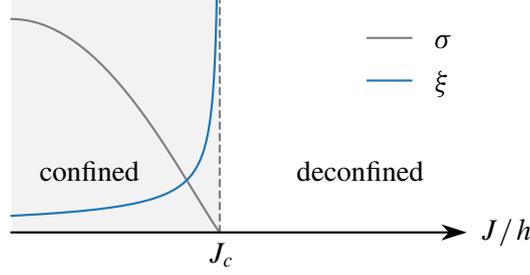


Figure 1.4: Schematic illustration of the confinement–deconfinement transition that occurs in the \mathbb{Z}_2 lattice gauge theory. When the magnetic field h is strong, the defects (violating the Gauss law constraint) are confined; the energy per unit length of separating the defects (the string tension), σ , is finite. This gives rise to finite confinement length scale ξ . Near the critical point, the string tension vanishes as $\sigma \sim (J - J_c)^{2\nu}$, while the confinement length diverges, $\xi \sim (J - J_c)^{-\nu}$. Since the transition belongs to the classical 3D Ising universality class, the critical exponents are well known, and $\nu \simeq 0.63$ [71, 72].

The \hat{A}_s operators are no longer conserved quantities, but the operators $\hat{B}_p = \prod_{i \in p} \hat{\sigma}_i^z$ are; they commute with each other, $[\hat{B}_p, \hat{B}_{p'}] = 0$, and with the Hamiltonian, $[\hat{H}_{\mathbb{Z}_2}, \hat{B}_p] = 0$. The conserved operators $\hat{G}_p \equiv \hat{B}_p$ generate a local (\mathbb{Z}_2) gauge transformation satisfying $\hat{G}_p^{-1} = \hat{G}_p$. The action of the \hat{G}_p on the spins is [similarly to Eq. (1.4.9)]

$$\hat{\sigma}_{qk}^x \rightarrow \hat{G}_p \hat{\sigma}_{qk}^x \hat{G}_p = e^{i\pi(\delta_{qp} + \delta_{kp})} \hat{\sigma}_{qk}^x, \quad (1.4.11)$$

where we have indexed the spins using the sites of the dual lattice ($\hat{\sigma}_{qk}$ lies on the link that bisects the line joining the centres of plaquettes q and k). That is, \hat{G}_p anticommutes with $\hat{\sigma}_{qk}^x$ if the support of the two operators overlaps. The gauge-invariant states satisfy a Gauss law constraint of the form $\hat{G}_p |\Psi\rangle = |\Psi\rangle$ for all plaquettes p . If the Hilbert space upon which (1.4.10) acts is restricted to such gauge-invariant states then we obtain the Ising gauge theory [68–70] (i.e., Gauss’ law is understood to be imposed on the states implicitly). It is then possible to introduce new gauge-invariant operators (Pauli matrices) $\hat{\tau}_s^x$ and $\hat{\tau}_s^z$ living on the sites of the lattice

$$\hat{\tau}_s^z = \hat{A}_s, \quad \hat{\sigma}_{ss'}^z = \hat{\tau}_s^x \hat{\tau}_{s'}^x, \quad (1.4.12)$$

in terms of which the Hamiltonian (1.4.10) becomes a two dimensional transverse field Ising model

$$\hat{H}_{\mathbb{Z}_2} = -h \sum_{\langle s, s' \rangle} \hat{\tau}_s^x \hat{\tau}_{s'}^x - J \sum_s \hat{\tau}_s^z. \quad (1.4.13)$$

Note that the duality relies on the implicit Gauss law constraint: When the transformation is written in the form $\hat{\tau}_s^x = \prod_{i \in \gamma(s)} \hat{\sigma}_i^z$, the path $\gamma(s)$ is only immaterial in the gauge-invariant subspace $\hat{B}_p |\Psi\rangle = |\Psi\rangle$.

The two dimensional Ising model (1.4.13) (whose Hilbert space is no longer restricted by a Gauss law constraint) exhibits two phases. For $h \gg J$, the Ising model is in its ferromagnetic phase* and $\langle \hat{\tau}_s^x \rangle > 0$. In the original $\hat{\sigma}_i^\alpha$ degrees of freedom, setting $J = 0$, the ground state is $\prod_j |\sigma_j^z = 1\rangle$ (which is explicitly gauge invariant). Excitations above this ground state correspond to flipping spins in the $\hat{\sigma}_j^z$ basis, with each spin flip costing an energy $2h$. Gauss' law enforces that these flipped spins must form closed loops on the lattice (since any endpoints would violate the constraint). This phase is *confining* in the sense that the energy required to separate two violations of the Gauss law constraint (i.e., plaquettes with $B_p = -1$) grows linearly with their separation relative to the ground state with no violations.

In the opposite limit, $J \gg h$, the Ising model is in its paramagnetic phase. If we were to ignore the Gauss law constraint, the ground state of the model would be $\prod_j |\sigma_j^x = 1\rangle$, which is clearly not gauge invariant. To obtain the gauge invariant ground state in this perturbative limit, one can project into the physical Hilbert space using the projection operator $\hat{P} = \prod_p \frac{1}{2}(\mathbb{1} + \hat{G}_p)$, which symmetrises over different gauge transformations[†]. Performing this projection leads us directly to the ground state of Kitaev's toric code model from the previous section, i.e., an equal-amplitude superposition of all possible closed loops of flipped spins generated by the \hat{G}_p . At the fine-tuned point $h = 0$, it costs strictly zero energy to separate two defects that violate the Gauss law constraint. For finite (but still sufficiently small) h , the defects interact via a screened potential of the form $V(\ell) \sim e^{-\ell/\xi_s}$, where ℓ is the separation of the defects [73], and the strength of the interaction and the screening length ξ_s depend on h and J . Deconfinement generally means that it costs a *finite* amount of energy to separate the defects by an infinite distance. At the critical value of h/J , the system undergoes a confinement–deconfinement phase transition; as the critical point is approached from the confining side, the string tension vanishes, and the corresponding confinement length diverges as shown schematically in Fig. 1.4.

1.4.3 Classical spin ice

Having introduced some of the central phenomenology relating to quantum spin liquids in the context of gapped \mathbb{Z}_2 theories, we now turn to looking at a robust mechanism that can produce spin liquid behaviour, both classical and quantum, in experiment. Namely, in the following, we will demonstrate how classical *Ising frustration* can give rise to classical spin liquid behaviour at low temperatures in dipolar spin ice. This discussion will set the stage for the introduction of *quantum* spin ice, where off-diagonal terms induce quantum fluctuations within the classical ground state

*Note that the ‘physical’ magnetic field acting on the $\hat{\sigma}_i^\alpha$ degrees of freedom maps onto the nearest-neighbour coupling in the dual Ising model.

[†]With periodic boundary conditions, $\prod_p \hat{G}_p = \mathbb{1}$, i.e., there are only 2^{N-1} independent gauge transformations. This allows the projector \hat{P} to be written in the form $\hat{P} = \frac{1}{2}\hat{S}(\mathbb{1} + \prod_p \hat{G}_p) = \hat{S}$, where \hat{S} symmetrises over independent gauge transformations.

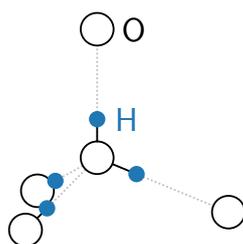


Figure 1.5: Arrangement of oxygen ions O^{2-} (white circles) and protons H^+ (smaller blue circles) satisfying the Bernal Fowler ice rules.

manifold, thereby generating a $U(1)$ quantum spin liquid.

Water ice

One of the hallmarks of frustrated magnetic systems is a ground state degeneracy that is exponentially large in the volume of the system. An important experimental consequence thereof is an extensive residual entropy as temperature is lowered to zero*. Historically, however, the first system that was shown in experiment to have a nonzero residual entropy was not a frustrated magnet. Instead, in the work of W. Giauque and collaborators [74, 75], it was shown that the much more familiar water ice (spin ice’s namesake) possesses such a residual entropy. In 1935, the origin of this result was explained quantitatively by Linus Pauling [76], who gave a simple counting argument for the residual entropy value that had been observed in experiment.

In the hexagonal phase of water ice, the O^{2-} ions form a tetrahedral structure, while the protons, H^+ , are located on the O–O bonds. The stability of the H_2O molecular structure leads to the so-called Bernal and Fowler ice rules [77], which describe where the protons should be placed within this tetrahedral structure of O^{2-} ions:

1. There should exist one proton per O–O bond,
2. In two of four bonds that emanate from each O^{2-} ion, the proton should be “close”, and in the remaining two it should be “far” from the Oxygen ion.

The second of these rules derives from the fact that the O–O bond length in hexagonal water ice significantly exceeds the short, covalent O–H bond in an H_2O molecule. A minimal structure that satisfies the above ice rules is shown in Fig. 1.5. Pauling argued in Ref. [76] that there are $\binom{4}{2} = 6$ ways of locally arranging the protons around each O^{2-} ion. However, each tetrahedron is not independent from its neighbours. If we consider a system that comprises N O^{2-} ions, then the structure contains $N_b = 2N$ O–O bonds. If the matching constraint (that each bond must contain

*The residual entropy appears experimentally when the system is cooled gradually to low temperatures. Ordering at low temperatures is generally expected, but the relaxation time scales at such low temperatures can be astronomically large, rendering the release of the residual entropy inaccessible.

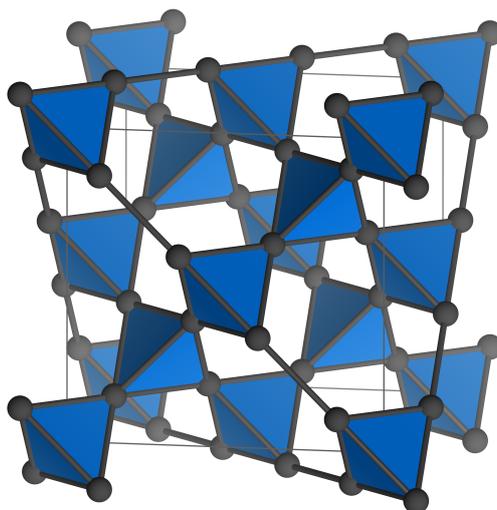


Figure 1.6: One cubic unit cell of the pyrochlore lattice. In spin ice, Ising spins live at the centres of the black spheres, between two corner-sharing tetrahedra, and are constrained to point along the local $\langle 111 \rangle$ direction (i.e., “in” or “out” of a given tetrahedron). The centres of the tetrahedra form a diamond lattice, which consists of two sublattices; “upwards” facing tetrahedra live on one sublattice, while “downwards” facing tetrahedra inhabit the complementary sublattice. The thin solid lines denote the conventional cubic unit cell.

one proton) is imposed on average, then one obtains $W = 6^N (1/2)^{N_b} = (3/2)^N$ configurations of the protons that are compatible with the Bernal Fowler ice rules. This value leads to a residual entropy $S = k_B \ln W = 0.806 \text{ cal/deg/mol}$, which is in close agreement with the experiments of Giauque [74] and more accurate analytical estimates that treat the constraints more carefully [78].

Spin ice

Spin ice compounds are part of the pyrochlore oxide family, which take the general form $A_2B_2O_7$. In such compounds, A is a rare-earth element such as Dy, Ho, while B is a non-magnetic ion such as Ti, Sn, Mo, Mn, etc. The ions A^{3+} and B^{4+} live on two interpenetrating pyrochlore sublattices. The corner-sharing tetrahedra that make up the pyrochlore lattice are shown in Fig. 1.6. The two canonical spin ice materials $\text{Dy}_2\text{Ti}_2\text{O}_7$ and $\text{Ho}_2\text{Ti}_2\text{O}_7$ are electronic insulators [79], and their magnetic ions Dy^{3+} and Ho^{3+} have total angular momentum $J = 15/2$ and $J = 8$, respectively. Strong crystal field effects in these materials lead to a ground state doublet that is separated from the first excited state by a gap of order 300 K. Within the ground state doublet, the total angular momentum \mathbf{J} is constrained to point along the local $\langle 111 \rangle$ direction, either “into” or “out” of the tetrahedra. In most cases* it is appropriate to project the full Hamiltonian into the ground space of the crystal field Hamiltonian (i.e., the ground state doublet on each site), and we are left with a

*The projection is appropriate as long as other energy scales in the problem do not become comparable to the crystal field splitting. Specifically, one requires temperatures $T \lesssim 10 \text{ K}$ and magnetic fields $B \lesssim 20 \text{ T}$ [80].

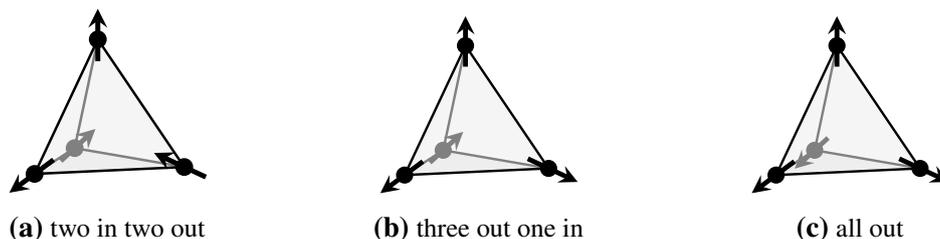


Figure 1.7: Symmetry-inequivalent configurations of the tetrahedra that make up the pyrochlore lattice. The two-in-two-out configuration in (a) does not host a charge, and corresponds to the ground state of an isolated tetrahedron. The configuration in (b) is an excited state that hosts a single charge (a monopole), $Q = 1$, while in the configuration in (c) the tetrahedron hosts a double charge, $Q = 2$.

model of effective, classical Ising spins $\sigma_i = -1, +1$ living on the sites of the pyrochlore lattice:

$$H_{\text{DSI}} = - \sum_{\langle i,j \rangle} J_{ij} (\mathbf{e}_i \cdot \mathbf{e}_j) \sigma_i \sigma_j + D \sum_{i < j} \frac{\mathbf{e}_i \cdot \mathbf{e}_j - 3(\mathbf{e}_i \cdot \mathbf{r}_{ij})(\mathbf{r}_{ij} \cdot \mathbf{e}_j)}{(r_{ij}/r_{\text{nn}})^3} \sigma_i \sigma_j. \quad (1.4.14)$$

The vectors \mathbf{e}_i point along the local $\langle 111 \rangle$ directions. The first term corresponds to the exchange interaction, and the second to long-ranged dipole–dipole interactions between the magnetic moments. Note that the sign of the effective interaction between the Ising spins σ_i is reversed, since $\mathbf{e}_i \cdot \mathbf{e}_j < 0$ for $i \neq j$.

Nearest neighbour spin ice

Nearest neighbour spin ice corresponds to the first term only in the Hamiltonian (1.4.14) (i.e., neglecting the dipolar contribution to the energy). Using the fact that $\mathbf{e}_i \cdot \mathbf{e}_j = -1/3$, we obtain

$$H_{\text{NN}} = J_{\text{eff}} \sum_{\langle i,j \rangle} \sigma_i \sigma_j, \quad (1.4.15)$$

where $J_{\text{eff}} = J/3^*$. The model (1.4.15) with antiferromagnetic $J_{\text{eff}} > 0$ on the pyrochlore lattice was first studied in the work of Anderson [82] in 1956 on cation ordering and antiferromagnetism in spinels. To find the ground state of the nearest neighbour spin ice Hamiltonian, we will begin by locally minimising the energy on a single tetrahedron. For an isolated tetrahedron, there are $2^4 = 16$ possible spin configurations, which may be classified, as in Fig. 1.7, according to the number of spins pointing “out” of the tetrahedron, and the number of spins pointing “in”. Specifically,

*In dipolar spin ice, the exchange interaction is antiferromagnetic. However, once the nearest neighbour contribution from the dipolar interaction is included, the sign of the effective nearest neighbour interaction between vector spins is made ferromagnetic. The $J > 0$ in this expression encompasses both the exchange interaction *and* the short-ranged contribution from the dipolar interaction. In real materials, $J_{\text{eff}} \simeq \frac{J}{3} + \frac{5D}{3}$, where $D > 0$ prevails over $J < 0$ in making $J_{\text{eff}} > 0$ [81].

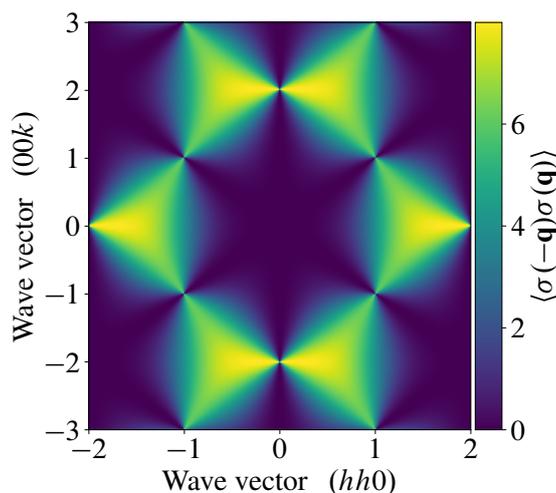


Figure 1.8: Pinch point structure in the static (equal time) structure factor $\langle \sigma(-\mathbf{q})\sigma(\mathbf{q}) \rangle$ of classical spin ice in the long-wavelength limit. The structure factor is plotted in the $[hhk]$ plane using the results of Ref. [83].

tetrahedra can be two-in-two-out ($\times 6$), three-in-one-out ($\times 4$), all out ($\times 1$), plus their spin-reversed counterparts. These different configurations of Ising spins can be conveniently parametrised by the “charge” on tetrahedron a as $Q_a = \pm \frac{1}{2} \sum_{i \in a} \sigma_i$. The sign depends on the sublattice that the tetrahedron belongs to, and ensures that, for example, a three-out-one-in tetrahedron (known as a magnetic monopole) will always obey $Q_a = 1$, irrespective of which sublattice it resides on. Similarly, a two-in-two-out tetrahedron will always be charge neutral, $Q_a = 0$. We may then note that the nearest neighbour Hamiltonian can be written solely in terms of the tetrahedral charges as

$$H_{\text{NN}} = 2J_{\text{eff}} \sum_a Q_a^2 - N_s J_{\text{eff}}, \quad (1.4.16)$$

from which we deduce that H_{NN} is both locally and *globally* minimised by $Q_a = 0$ on every tetrahedron. Although this constraint is trivial in the charge language, it is nontrivial when imposed on the underlying Ising spin degrees of freedom. Indeed, $Q_a = 0$ is equivalent to requiring that all tetrahedra satisfy the two-in-two-out constraint. The equivalence with water ice is now clear: If in water ice we associate an Ising variable that labels whether the proton is in the “near” or the “far” position, then the Bernal and Fowler ice rules are equivalent to the two-in-two-out constraint in spin ice. The nearest neighbour Ising Hamiltonian (1.4.16) therefore gives rise to a finite zero point entropy given approximately by the Pauling [76] result $S_0/k_B N_s \simeq (1/2) \log(3/2)$.

Writing the monopole-free constraint in terms of the original vector spin degrees of freedom, $\mathbf{S}_i = \sigma_i \mathbf{e}_i$, we find that the two-in-two-out constraint is equivalent to $\sum_{i \in a} \mathbf{S}_i = 0$ on each tetrahedron a . This constraint on the vector spins can be viewed as a lattice regularisation of the divergence: $\sum_{i \in a} \mathbf{S}_i \equiv (\nabla \cdot \mathbf{S})_a = 0$. Therefore, upon coarse graining, the spin vector field $\mathbf{S}(\mathbf{r})$,

which is proportional to the physical magnetisation, satisfies the solenoidal constraint $\nabla \cdot \mathbf{S} = 0$, analogous to the Maxwell equation $\nabla \cdot \mathbf{B} = 0$ for the magnetic field in classical electromagnetism. The divergence-free condition can be satisfied by introducing an emergent gauge field $\mathbf{A}(\mathbf{r})$, such that $\mathbf{S}(\mathbf{r}) = \nabla \times \mathbf{A}(\mathbf{r})$. Endowing the continuum theory with a Gaussian weight functional*, $\rho[\mathbf{A}(\mathbf{r})] \propto \exp(-\frac{K}{2} \int d\mathbf{r} (\nabla \times \mathbf{A})^2)$ [83], leads to a Maxwell-like action, with dipolar correlations between spins

$$\langle \sigma_i(\mathbf{r}) \sigma_j(0) \rangle \propto \frac{3r_i r_j - r^2 \delta_{ij}}{r^5}. \quad (1.4.17)$$

In reciprocal space, these correlations give rise to pinch point singularities (as depicted in Fig. 1.8), which have been observed in neutron scattering experiments on spin ice compounds [85, 86], and constitute compelling experimental evidence for the existence of emergent gauge fields and the associated Gauss law constraint.

Including the dipolar interaction

Although many of the properties of spin ice, such as the extensive zero temperature entropy and the lack of any finite temperature phase transition, are explained by the simple nearest neighbour model (1.4.16), the full Hamiltonian includes also the long-ranged dipolar interaction, whose magnitude is comparable to the exchange interaction. By virtue of projective equivalence [87], the dipolar interactions can be accounted for effectively by a long-range, pairwise interaction between the tetrahedral charges Q_α

$$H_c = 2J_{\text{eff}} \sum_a Q_a^2 + E_{\text{nn}} \sum_{a < b} \frac{Q_a Q_b}{r_{ab}}, \quad (1.4.18)$$

where $r_{ab} = |\mathbf{r}_a - \mathbf{r}_b|/r_{\text{nn}}$ is the distance between the centres of the tetrahedra a and b in units of the distance between neighbouring tetrahedra, and E_{nn} is the Coulomb energy of two adjacent unit charges. Such a description can also be justified by writing each dipole as the sum of two fictitious monopoles [88], in which case (1.4.18) is known as the dumbbell model. The first term accounts for both the exchange interaction and a short-ranged contribution from the dipolar term. The latter accounts for the long-ranged nature of the dipole interaction. Projective equivalence ensures that (1.4.18) is quantitatively accurate up to quadrupolar corrections. When E_{nn} is nonzero, the ground state(s) are determined by a competition between an ordered arrangement of charges on the diamond lattice and the charge vacuum. For E_{nn} sufficiently small, the self-energy cost of the charges prevails, and the $Q_\alpha = 0$ configurations remain *exact* ground states of the approximate (1.4.18).

The monopole description (1.4.18) provides an excellent approximation to both the equilibrium behaviour [89, 90] and dynamics [91, 92] of classical spin ice. However, one area where the

*Entropically, the system favours spin configurations that form lots of short loops, which upon coarse graining map to small values of $\mathbf{S}(\mathbf{r})$ [83, 84].

dumbbell model (at least in theory) falls short is the absence of a low temperature phase transition. Equation (1.4.18) predicts that the nearest-neighbour spin ice ground states, characterised by $Q_a = 0$, are exactly degenerate. Reinstating the full dipolar interaction splits the degeneracy of such states, which leads to a first order phase transition at low temperatures [93, 94] (approximately 180 mK for $\text{Dy}_2\text{Ti}_2\text{O}_7$ [80]), selecting a particular ordered spin state from the ground state manifold. In order to observe such a transition numerically (using the Monte Carlo methods introduced later in Sec. 4.2), one must implement loop updates [94], allowing the system to efficiently pass between states that satisfy the ice rules at low temperatures. However, the dynamics of classical spin ice, both in and out of equilibrium, is described quantitatively by “single spin flip” dynamics [84]. The equilibration time scale under such dynamics diverges exponentially at low temperatures [91, 92], rendering the ordering transition experimentally inaccessible in practice [95–97].

1.4.4 Quantum spin ice and U(1) lattice gauge theory

The canonical model of *quantum* spin ice consists of the nearest-neighbour classical spin ice Hamiltonian and transverse (exchange) terms, which induce quantum fluctuations. Specifically, one typically considers a local XXZ model on the pyrochlore lattice [98]:

$$\hat{H}_{\text{QSI}} = J_z \sum_{\langle ij \rangle} \hat{S}_i^z \hat{S}_j^z - J_\perp \sum_{\langle ij \rangle} \left(\hat{S}_i^+ \hat{S}_j^- + \hat{S}_i^- \hat{S}_j^+ \right), \quad (1.4.19)$$

where the transverse term $\propto J_\perp$ is assumed small, $|J_\perp| \ll J_z$. As discussed in the previous section, if we neglect the transverse terms entirely, the ground states of the model are \hat{S}_i^z tensor product states that satisfy the ice rules, with $\sum_{i \in t} \hat{S}_i^z = 0$ on each tetrahedron t . Turning on the transverse terms perturbatively, one may then use degenerate perturbation theory to study how the extensive degeneracy of the ground state sector is split. The degeneracy is first split at *third* order* in J_\perp , where ring exchange processes give rise to the effective Hamiltonian [100]

$$\hat{H}_{\text{ring}} \propto \frac{J_\perp^3}{J_z^2} \sum_{h \in \{\square\}} \hat{S}_{h,1}^+ \hat{S}_{h,2}^- \hat{S}_{h,3}^+ \hat{S}_{h,4}^- \hat{S}_{h,5}^+ \hat{S}_{h,6}^- + \text{H.c.}, \quad (1.4.20)$$

which acts on the Hilbert space spanned by the ground states of classical spin ice. Equivalently, we can view this from the perspective of a Gauss law constraint that selects the physical Hilbert space: $\sum_{i \in t} \hat{S}_i^z |\Psi\rangle = 0$. The Gauss law constraint acts as the generator of infinitesimal gauge transformations of the form $e^{i \sum_t \alpha_t (\sum_{i \in t} \hat{S}_i^z)}$, i.e., rotations about the local z axis of all spins belonging to each tetrahedron. The Hilbert space corresponding to the ground state sector is then spanned by gauge invariant states, which satisfy $e^{i \sum_t \alpha_t (\sum_{i \in t} \hat{S}_i^z)} |\Psi\rangle = |\Psi\rangle$.

*At second order in J_\perp , creation and annihilation of virtual, constraint-violating defects leads to a constant shift in all energy levels [99].

To write the ring exchange Hamiltonian (1.4.20) in terms of a U(1) lattice gauge theory, the spin degrees of freedom can be written in terms of canonically conjugate rotor variables, \hat{n}_i and $\hat{\phi}_i$, where $\hat{S}_i^z = \hat{n}_i - \frac{1}{2}$ and \hat{S}_i^\pm act as raising/lowering operators $\hat{S}_i^\pm = e^{\pm i\hat{\phi}_i}$ [99]. The hard constraint $n_i = 0, 1$ can be implemented at the level of the Hamiltonian by ensuring that the system incurs a finite cost if the constraint is violated. In the rotor variables, the ring exchange Hamiltonian becomes

$$\hat{H} = \frac{U}{2} \sum_{\langle \mathbf{r}\mathbf{r}' \rangle} [(\hat{n}_{\mathbf{r}\mathbf{r}'} - \frac{1}{2})^2 - \frac{1}{4}] - K \sum_{h \in \{\diamond\}} \cos\left(\hat{\phi}_{h,1} - \hat{\phi}_{h,2} + \hat{\phi}_{h,3} - \hat{\phi}_{h,4} + \hat{\phi}_{h,5} - \hat{\phi}_{h,6}\right), \quad (1.4.21)$$

where the vector indices \mathbf{r} label the sites of the diamond lattice (i.e., the centres of the tetrahedra). States with $n_{\mathbf{r}\mathbf{r}'} \neq 0, 1$ are penalised by an energy cost $\propto U$; only in the limit $U/K \rightarrow \infty$ is the hard constraint $n_i = 0, 1$ enforced exactly. The introduction of (directed) link variables $\hat{e}_{\mathbf{r}\mathbf{r}'} = \pm(\hat{n}_{\mathbf{r}\mathbf{r}'} - \frac{1}{2})$ and $\hat{a}_{\mathbf{r}\mathbf{r}'} = \pm\hat{\phi}_{\mathbf{r}\mathbf{r}'}$ (where the sign is positive for \mathbf{r} belonging to the A sublattice, which we take to be composed of ‘‘upwards’’ facing tetrahedra) makes the correspondence with compact U(1) lattice gauge theory explicit*

$$\hat{H} = \frac{U}{2} \sum_{\langle \mathbf{r}\mathbf{r}' \rangle} \hat{e}_{\mathbf{r}\mathbf{r}'}^2 - K \sum_{h \in \{\diamond\}} \cos\left(\sum_{\mathbf{r}\mathbf{r}' \in h} \hat{a}_{\mathbf{r}\mathbf{r}'}\right), \quad (1.4.22)$$

with the lattice curl $(\nabla \times \hat{a})_h \equiv \sum_{\mathbf{r}\mathbf{r}' \in h} \hat{a}_{\mathbf{r}\mathbf{r}'}$. The canonically conjugate variables $\hat{a}_{\mathbf{r}\mathbf{r}'}$ and $\hat{e}_{\mathbf{r}\mathbf{r}'}$ play the role of the vector and electric potentials, respectively. Violations of Gauss’ law correspond to electric charges, and in the ground state sector $(\nabla \cdot \hat{e})_{\mathbf{r}} \equiv \sum_{\mathbf{r}'} \hat{e}_{\mathbf{r}\mathbf{r}'} = 0$. In $d = 3$ spatial dimensions, the unfrustrated analogue of (1.4.22) exhibits two phases [101]; for $U/K \ll 1$ the theory exhibits a deconfined phase, while in the opposite limit the theory is confining. In the deconfined phase[†], the compact nature of the gauge theory is not important, the cosine can be expanded to quadratic order, and one arrives at the familiar Maxwell Hamiltonian of (noncompact) U(1) quantum electrodynamics (QED)

$$\hat{H} = \frac{U}{2} \sum_{\langle \mathbf{r}\mathbf{r}' \rangle} \hat{e}_{\mathbf{r}\mathbf{r}'}^2 + \frac{K}{2} \sum_{h \in \{\diamond\}} (\nabla \times \hat{a})_h^2. \quad (1.4.23)$$

By analogy with QED, the Hamiltonian (1.4.23) exhibits gapless, linearly dispersing degrees of freedom at low energies: the photon. Gauge invariance implies that the photon has two polarisations, both of which are transverse to its momentum. Electric charges (otherwise known as spinons),

*In the limit $U/K \rightarrow \infty$, the Hamiltonian enforces half-integer values of the electric field, $e_{\mathbf{r}\mathbf{r}'} = \pm\frac{1}{2}$. The gauge theory is therefore referred to as *frustrated* [99] or *odd* [73, 98], in contrast to unfrustrated gauge theories in which the strong coupling limit $U/K \rightarrow \infty$ enforces the trivial vacuum $e_{\mathbf{r}\mathbf{r}'} = 0$.

[†]Monte Carlo simulations suggest that, in the frustrated/odd gauge theory, the deconfined phase survives all the way to $U/K \rightarrow \infty$ [102, 103].

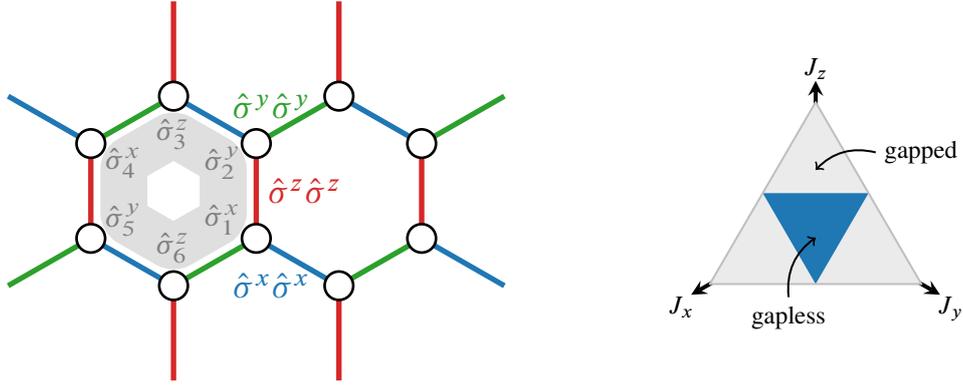


Figure 1.9: **Left:** Illustration of the anisotropic interactions and the conserved plaquette operator \hat{W}_p in the Kitaev model. Along the blue (green, red) bonds, the spins interact via interactions $\hat{\sigma}_i^\gamma \hat{\sigma}_j^\gamma$ with $\gamma = x$ (y, z). **Right:** Phase diagram of the Kitaev model in the plane $J_x + J_y + J_z = 1$ (perpendicular to $[111]$).

which correspond to violations of the Gauss law constraint, $(\nabla \cdot \hat{e})_r = 0$, have an energy gap of order J_z , and interact via an emergent Coulomb interaction [99, 104], mediated by the photon. In addition to the energetically costly spinons and the gapless “photon”, the noncompact nature of the U(1) gauge theory permits quantised magnetic charge. These magnetic charges (otherwise known as visons) are constructed from states belonging to the spin ice manifold, and possess a finite gap set by J_\perp^3/J_z^2 [99], intermediate between the energies of the photon and the spinon excitations. Pairs of vison excitations also interact via a long-ranged $1/r$ Coulomb potential.

1.4.5 Kitaev’s honeycomb model

In the last twenty years, two exactly solvable models stand out as being particularly influential in aiding our understanding of quantum spin liquids and their associated phenomena, both of which are due to Kitaev. The first is the toric code [46], introduced in Sec. 1.4.1, which illustrates the pertinent physics associated with topological order in gapped \mathbb{Z}_2 QSLs. The second is Kitaev’s honeycomb model [47], which, as we will present here, possesses the trinity of exact solvability, experimental relevance, and a rich phase diagram that displays a variety of gapped and gapless quantum spin liquid ground states.

The Kitaev Hamiltonian is exactly solvable on any tricoordinated lattice [50], but in the following we will restrict our attention to the two dimensional honeycomb lattice (see Fig. 1.9), which was the subject of Kitaev’s seminal work in 2006 [47]. The Hamiltonian is then

$$\hat{H}_{\text{Kitaev}} = - \sum_{\gamma=x,y,z} \sum_{\langle j,k \rangle_\gamma} J^\gamma \hat{\sigma}_j^\gamma \hat{\sigma}_k^\gamma, \quad (1.4.24)$$

where the spin-1/2 degrees of freedom $\hat{\sigma}_j^\gamma$ live on the sites of the honeycomb lattice, and γ labels the

bond direction. The nearest-neighbour interaction between the spins is highly anisotropic, with a bond-direction-dependent easy axis* (the defining feature of so-called compass models [111]). This anisotropy, which is at the core of the model's highly unusual properties, arises in some transition metal compounds due to strong spin-orbit coupling (for reviews of candidate spin liquid compounds, including Kitaev materials that possess such anisotropy, see, e.g., [50, 51, 56, 112–116]). In the vicinity of the isotropic point $J^x = J^y = J^z$, the ground state of the Hamiltonian is a *gapless* spin liquid. Conversely, when one of the coupling constants is made significantly larger than the other two, the lowest order effective Hamiltonian is precisely that of the toric code [47], and the system is correspondingly gapped, as shown in the phase diagram in Fig. 1.9. The Hamiltonian (1.4.24) possesses an extensive number of conserved quantities, \hat{W}_p , one for each plaquette p , defined by

$$\hat{W}_p = \prod_{\langle i,j \rangle_\gamma \in p} \hat{\sigma}_i^\gamma \hat{\sigma}_j^\gamma = \hat{\sigma}_1^x \hat{\sigma}_2^y \hat{\sigma}_3^z \hat{\sigma}_4^x \hat{\sigma}_5^y \hat{\sigma}_6^z, \quad (1.4.25)$$

where the numerical subscripts index the spins belonging to the plaquette p in an anti-clockwise direction, as shown in Fig. 1.9. Since $\hat{W}_p^2 = \mathbb{1}$, the eigenvalues of the plaquette operator must be ± 1 . Although the number of such operators is extensive, specifying the eigenvalue of each plaquette operator does not uniquely determine the spins. Indeed, there are two spins per plaquette, but only one constraint, and hence there remain an extensive number of degrees of freedom unspecified.

In Kitaev's original exact solution of the model, the spins are represented (with some redundancy) as Majorana fermions. Specifically, one introduces *four* Majorana operators per site, \hat{a}_j^α ($\alpha = x, y, z$) and \hat{c}_j , allowing the spins to be represented as [47, 57]

$$\hat{\sigma}_j^\alpha = i \hat{a}_j^\alpha \hat{c}_j. \quad (1.4.26)$$

The spins satisfy the canonical angular momentum commutation relations as long as $\hat{a}_j^x \hat{a}_j^y \hat{a}_j^z \hat{c}_j = \mathbb{1}$. Therefore, the physical subspace is obtained via the application of the projection operator $\hat{P}_j = \frac{1}{2}(\mathbb{1} + \hat{a}_j^x \hat{a}_j^y \hat{a}_j^z \hat{c}_j)$ on each site j . In the Majorana language, the bond operators become $\hat{\sigma}_i^\gamma \hat{\sigma}_i^\gamma = -(i \hat{a}_i^\gamma \hat{a}_i^\gamma) i \hat{c}_i \hat{c}_j$. Further defining $\hat{u}_{ij} = i \hat{a}_i^{\gamma(i,j)} \hat{a}_j^{\gamma(i,j)}$, and $\hat{A}_{ij} = 2J^{\gamma(i,j)} \hat{u}_{ij}$, the Hamiltonian reduces to

$$\hat{H}_{\text{Kitaev}} = \frac{i}{4} \sum_{\langle i,j \rangle} \hat{A}_{i,j} \hat{c}_i \hat{c}_j. \quad (1.4.27)$$

Since the bilinear \hat{u}_{ij} commute with one another and with \hat{H}_{Kitaev} , they can be replaced by their eigenvalues, ± 1 , in Eq. (1.4.27). The problem (in the enlarged Hilbert space) is hence reduced to Majorana fermions hopping through a \mathbb{Z}_2 gauge field, as specified by the eigenvalues of the \hat{u}_{ij} .

*In more realistic Hamiltonians, it is common to include some residual isotropic Heisenberg interactions between nearest neighbours (the Kitaev-Heisenberg model [105]), and/or symmetric off-diagonal exchange interactions, in which case the Hamiltonian is known as the $JK\Gamma$ [106–108]/ $K\text{-}\Gamma$ [109, 110] model.

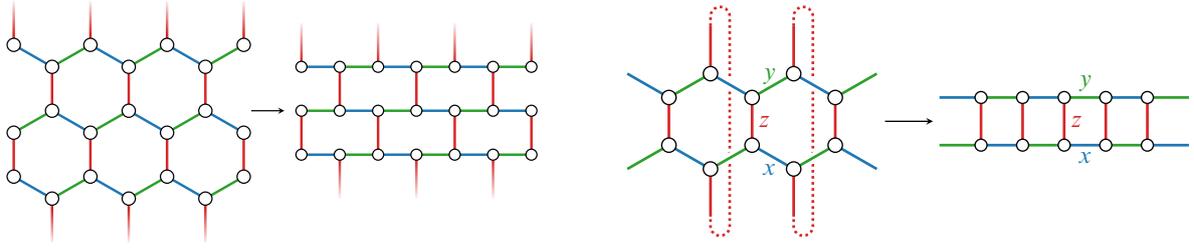


Figure 1.10: **Left:** The honeycomb lattice can alternatively be viewed as a “brick wall” lattice in which the spins are arranged on the sites of a square lattice. Rows of spins are connected in the vertical direction on either odd or even sites depending on the parity of the row. **Right:** The Kitaev honeycomb model with one row of hexagons and periodic boundary conditions can, equivalently, be viewed as a model that resides on a square lattice two-leg ladder with a unit cell of length two.

Choosing $u_{ij} = \pm 1$ is equivalent to “gauge fixing” [50], while projection into the physical subspace $\hat{P}_j |\Psi\rangle = |\Psi\rangle$ (analogous to Gauss’ law in conventional electromagnetism) may be considered as symmetrisation over different gauges. Fortunately, a theorem due to Lieb [117] can be used to prove that, from the exponentially large number of flux sectors, the one containing the ground state of the model is the one that contains no \mathbb{Z}_2 flux, i.e., $W_p = 1$ for all plaquettes p . After gauge fixing, the Hamiltonian (1.4.27) in the flux free sector can be written using $u_{ij} = 1$ (-1) if i belongs to the A (B) sublattice. The fermionic Hamiltonian can then be explicitly diagonalised to obtain the dispersion relation $\epsilon_{\mathbf{k}} = 2|J^x e^{i\mathbf{k}\cdot\mathbf{n}_1} + J^y e^{i\mathbf{k}\cdot\mathbf{n}_2} + J^z|$ [with $\mathbf{n}_{1,2} = (\pm\sqrt{3}/2, 3/2)$], which, in the gapless phase, hosts two zero-energy Dirac points [47, 57].

Once again we observe that the system harbours an emergent gauge field. Here, outside of the ground state sector, the \mathbb{Z}_2 valued gauge field acts as a static source of π flux (i.e., half of the flux quantum) for the Majorana fermions. That is, when a fermion encircles an odd number of hexagonal plaquettes with $W_p = -1$, it acquires an Aharonov–Bohm phase of $e^{i\pi} = -1$. The other defining features of quantum spin liquids introduced previously are also present in the Kitaev model. For instance, it follows from (1.4.26) that spins fractionalise into a pair of nearest neighbour, static π fluxes and a Majorana fermion [118]. Moreover, in both the gapped and gapless phases, the ground state exhibits long-range entanglement, generating a subleading, topological contribution to the entanglement entropy [119], and the ground state degeneracy depends on the topology of the manifold on which the model resides [120–122].

From honeycomb to ladders

The Kitaev Hamiltonian (1.4.24) can alternatively be defined on a two-leg spin ladder [67, 123, 124]. We will introduce this model briefly here, since it possesses a number of key conceptual similarities with compass model that we will present later in Chapter 5. The conclusions that we draw in Chapter 5 can then be applied *mutatis mutandis* to the Kitaev ladder model. The honeycomb

lattice can equivalently be viewed as “brick wall” lattice, as shown in Fig. 1.10. The spins are then labelled according to their column (j) and row (ℓ) indices, and the Kitaev Hamiltonian (1.4.24) becomes [123]

$$\hat{H}_{\text{Kitaev}} = - \sum_{j+\ell=\text{even}} \left(J_x \hat{\sigma}_{j,\ell}^x \hat{\sigma}_{j+1,\ell}^x + J_y \hat{\sigma}_{j-1,\ell}^y \hat{\sigma}_{j,\ell}^y + J_z \hat{\sigma}_{j,\ell}^z \hat{\sigma}_{j,\ell+1}^z \right) \quad (1.4.28)$$

Putting the Kitaev Hamiltonian on a single “row of bricks” (with periodic boundary conditions) leads to the ladder generalisation of the Kitaev model, which is shown schematically in Fig. 1.10. Following Ref. [123], the spins in (1.4.28) can then be fermionised via a Jordan–Wigner transformation along the rows of the lattice to complex fermionic operators $\hat{a}_{j,\ell}, \hat{a}_{j,\ell}^\dagger$, defined explicitly by the relations

$$\hat{\sigma}_{j,\ell}^+ = \hat{a}_{j,\ell}^\dagger \exp \left(i\pi \sum_{i,k<\ell} \hat{n}_{i,k} + i\pi \sum_{i<j} \hat{n}_{i,\ell} \right). \quad (1.4.29)$$

Unlike Kitaev’s original solution, this fermionic representation involves *no redundancy*, i.e., there is no Gauss law constraint that selects physical (gauge invariant) states from an enlarged Hilbert space. In addition, this mapping offers additional utility in that it may be used to define a string order parameter, which becomes local within a dual description, allowing the phase transition between the gapped and gapless phases to be characterised within Landau’s theory of continuous phase transitions [123, 125]. The complex fermions (\hat{a}, \hat{a}^\dagger) can then be transformed to Majorana fermions (\hat{c}, \hat{d}) with the convention [123]

$$\hat{d}_{j\ell} = \begin{cases} \hat{a}_{j\ell}^\dagger + \hat{a}_{j\ell} & \text{if } j + \ell = \text{even,} \\ i(\hat{a}_{j\ell}^\dagger - \hat{a}_{j\ell}) & \text{if } j + \ell = \text{odd,} \end{cases} \quad (1.4.30)$$

$$\hat{c}_{j\ell} = \begin{cases} i(\hat{a}_{j\ell}^\dagger - \hat{a}_{j\ell}) & \text{if } j + \ell = \text{even,} \\ \hat{a}_{j\ell}^\dagger + \hat{a}_{j\ell} & \text{if } j + \ell = \text{odd.} \end{cases} \quad (1.4.31)$$

These transformations lead to a representation in which the conserved quantities of the Kitaev Hamiltonian are explicit*

$$\hat{H}_{\text{Kitaev}} = i \sum_{j+\ell=\text{even}} \left(J_x \hat{c}_{j,\ell} \hat{c}_{j+1,\ell} - J_y \hat{c}_{j-1,\ell} \hat{c}_{j,\ell} + J_z (i \hat{d}_{j,\ell} \hat{d}_{j,\ell+1}) \hat{c}_{j,\ell} \hat{c}_{j,\ell+1} \right). \quad (1.4.32)$$

*Note that periodic boundary conditions are imposed in the row index ℓ such that $\ell + 1$ is understood to mean $\ell + 1 \bmod 2$.

Specifically, the operators $\hat{\eta}_{j\ell} = i\hat{d}_{j,\ell}\hat{d}_{j,\ell+1}$ (of which there are L) commute with the Hamiltonian and one another. Hence, they correspond to good quantum numbers, $\eta_{j\ell} = \pm 1$, and the Hamiltonian can be block diagonalised using a basis in which the $\hat{\eta}_{j\ell}$ are simultaneously diagonalised. The conserved quantities are analogous to the \mathbb{Z}_2 fluxes \hat{W}_p in the honeycomb model [see Eq. (1.4.25)]. While the ground state is given by a configuration of conserved quantities that are translationally invariant with respect to the unit cell [67, 117, 123, 124], within a *typical* block, the dynamics is that of Majorana fermions in the presence of a binary random field, $\pm J_z$. To describe the behaviour of the fermions in the presence of such emergent randomness, we are led to consider the phenomenon of Anderson localisation.

1.5 Localisation phenomena

1.5.1 Anderson localisation

Anderson localisation is a surprising and inherently quantum mechanical phenomenon that arises in random media owing to the wavelike nature of particles at microscopic scales. To illustrate some of the main counter-intuitive implications, imagine a single particle of mass m in one spatial dimension subjected to the Hamiltonian

$$H = \frac{p^2}{2m} + V(x), \quad (1.5.1)$$

where $V(x)$ is a random potential (the details of which will remain unspecified for now). If we first think classically, there are two relevant regimes. First, if the kinetic energy of the particle significantly exceeds the typical fluctuations of the potential, the particle is free to move above the potential landscape essentially unimpeded. On the other hand, if the kinetic energy of the particle is small compared to the typical fluctuations of the potential, it will remain stuck and oscillate back and forth, i.e., classically localised, within some region of space set by energy conservation.

When treating the problem quantum mechanically, it is well known that a change in the potential leads to reflection and transmission of the wave function. The relative magnitude of reflection and transmission depends on the height of the barrier, but *some* probability will be transmitted through the barrier, even when the barrier's height exceeds the kinetic energy of the particle, i.e., quantum tunnelling. Naïvely one may therefore expect that if each 'bump' in the potential leads to reflection and transmission, the probability density should be governed by a random walk and therefore exhibit diffusive propagation characterised by $\langle x^2 \rangle \sim t$. This intuition turns out to be incorrect, at least in spatial dimensions $d \leq 2$, or in the presence of sufficiently strong disorder.

In Anderson's seminal work [126], it was shown that – contrary to the heuristic arguments above – the role of *interference* between the various incident and reflected waves cannot be disregarded and causes the particle to remain stuck in the vicinity of its initial position. At long times, if all eigenstates are localised, the average dynamics of the particle ceases, $\langle x^2 \rangle \sim \text{const.}$, and the wave function of the particle (asymptotically) decays exponentially in space

$$|\psi(\mathbf{r})|^2 \sim \exp\left(-\frac{|\mathbf{r}|}{\xi}\right), \quad (1.5.2)$$

where the position \mathbf{r} is measured with respect to the initial position of the particle. The length scale ξ that controls the exponential decay is the localisation length. Since interference lies at the heart of localisation phenomena, localisation effects are also relevant in other disordered systems described by (classical) wave equations. For instance, Anderson localisation can occur in the random scattering of light [127, 128], matter waves [129, 130], and even sound waves [131]. In

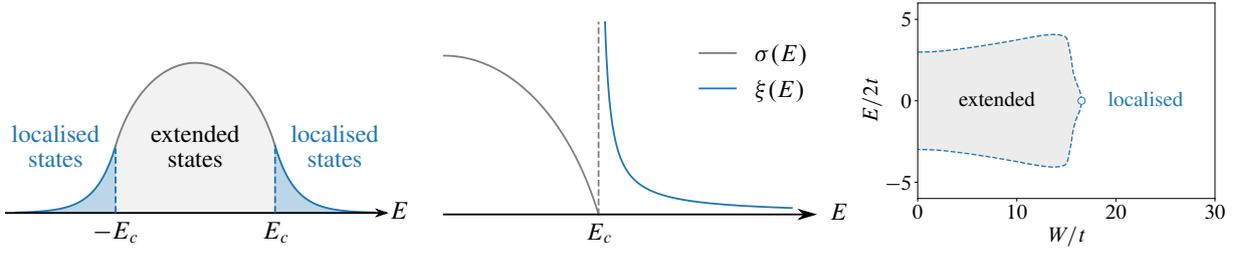


Figure 1.11: **Left:** Schematic depiction of the density of states and mobility edges in the three-dimensional Anderson model. The states in the centre of the band, with energies $|E| < E_c$, are extended, while those in the band edges, $|E| > E_c$, are localised. **Centre:** In the vicinity of E_c we plot the localisation length ξ and the conductivity σ resolved by energy. On the metallic side of the transition, $|E| < E_c$, the conductivity vanishes as $\sigma(E) \sim (E - E_c)^s$, while on the insulating side, $|E| > E_c$, the localisation length diverges as $\xi(E) \sim (E_c - E)^{-\nu}$. **Right:** Mobility edges as a function of disorder for the simple cubic lattice, inspired by Ref. [132]. As disorder is increased from zero, the mobility edges initially move outwards, before they eventually turn around and move towards the centre of the band. Above a critical disorder strength, W_c , denoted by the white circle, all eigenstates become exponentially localised.

the following, we will discuss some of the most common characterisations of localisation that complement the intuitive “absence of diffusion” definition given above.

If the particle is initially located on a single lattice site, an alternative but related notion of localisation is to think about the probability that the particle returns to its initial position (the return probability). If the initial state has nonzero overlap with localised states then some fraction of the initial particle density will not diffuse away from its initial position. Introducing the Green’s function

$$\mathcal{G}(\mathbf{r}, \mathbf{r}'; t) = \sum_{\alpha} \psi_{\alpha}(\mathbf{r}) \bar{\psi}_{\alpha}(\mathbf{r}') e^{-iE_{\alpha}t}, \quad (1.5.3)$$

where $\psi_{\alpha}(\mathbf{r}) = \langle \mathbf{r} | \alpha \rangle$ are the eigenfunctions of the system in real space, the return probability is given by $|\mathcal{G}_{00}(t)|^2$. One can then show that for sufficiently large times [133]

$$\lim_{t \rightarrow \infty} |\mathcal{G}_{00}(t)|^2 = \sum_{\alpha} |\langle 0 | \alpha \rangle|^4. \quad (1.5.4)$$

If the right hand side does not vanish, then some fraction of the probability will remain stuck in the vicinity of the origin, defined by the particle’s initial position. This expression also leads us to another important quantity when considering localised systems: the (inverse) participation ratio (IPR). The IPR also allows for the definition of so-called multifractal states, which are extended but non-ergodic. Such multifractal states, which are characterised by rare peaks and algebraic tails, occur in a variety of contexts, perhaps most notably at the Anderson and quantum Hall plateau

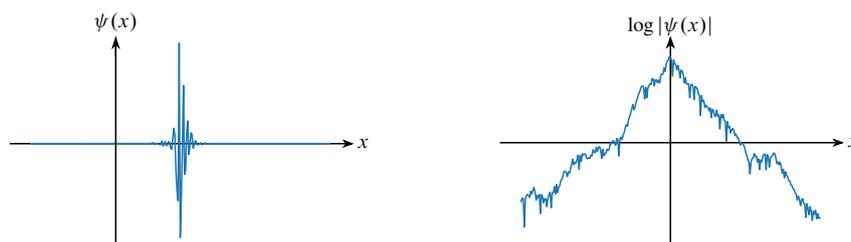


Figure 1.12: **Left:** Typical mid-spectrum state from the one-dimensional Anderson model on a lattice. The state is localised over a region of size $\sim \xi$. **Right:** The logarithm of the same mid-spectrum state, highlighting the state’s exponential tails.

transitions [134–137] and in hierarchical structures [138, 139]. The IPR is defined as

$$\text{IPR}_q^\alpha = \sum_{i=1}^N |\langle i|\alpha\rangle|^{2q}, \quad (1.5.5)$$

where we have assumed the normalisation $\text{IPR}_1^\alpha = 1$. One common way to distinguish between extended and localised states is to look at the fourth power ($q = 2$) of the wave function. In the case of perfect extended Bloch waves, $\text{IPR}_2^\alpha = 1/N$. For random extended states, the prefactor is modified, but one generally expects the $\propto N^{-1}$ scaling to hold [133]. On the contrary, for perfectly localised states, the wave functions are Kronecker delta functions $\psi_i^\alpha = \delta_i^\alpha$, and the summation gives $\text{IPR}_2^\alpha = 1$. Generalising to exponentially localised states with appreciable support over a region $\sim \xi$, the IPR does not scale with the volume of the system and assumes the value $\text{IPR}_2^\alpha \sim \xi^{-d}$. Thence, in the thermodynamic limit, whether IPR_2 vanishes or not indicates whether the system possess localised or extended states at the energy being examined. More generally, for a system of volume V , $\overline{\text{IPR}}_q \sim V^{-\tau_q}$, where the overline corresponds to disorder averaging. The moments $\{\tau_q\}$ are commonly written in terms of the so-called multifractal exponents $\{D_q\}$

$$\tau_q = D_q(q - 1). \quad (1.5.6)$$

Localised systems have $D_q = 0$, extended systems have $D_q = 1$, while for multifractal states D_q depends nontrivially on the power q .

One dimensional systems

After the original work of Anderson [126], it was conjectured in the early 1960s that in one dimensional systems, *all* single particle wave functions are exponentially localised in the presence of spatially uncorrelated disorder, irrespective of its strength [140]. This conjecture was later given a more mathematically rigorous footing by Goldshtein and collaborators [141]. In 1979, the “gang of four”, E. Abrahams, P. W. Anderson, D. C. Liccardello, and T. V. Ramakrishnan, published their

one parameter scaling theory of Anderson localisation [142], which we will soon discuss. Although the approach is not rigorous, it has played a significant role in the field. In particular, it was the first time that the role of dimensionality was elucidated, establishing that the critical dimension for the existence of transport is $d = 2$, providing further evidence that all states in $d = 1$ are localised.

In one dimensional and quasi-one dimensional systems, the localisation length can also be defined through the Lyapunov exponents of products of random matrices. This definition is particularly convenient for the purposes of numerical calculations. Consider, for example, a strictly one-dimensional tight binding model

$$\hat{H}_{\text{Anderson}} = -t \sum_{n=1}^{L-1} \left(c_n^\dagger c_{n+1} + c_{n+1}^\dagger c_n \right) + \sum_{n=1}^L V_n c_n^\dagger c_n, \quad (1.5.7)$$

where we impose open boundary conditions. When the chain is clean, $P(V_n) = \delta(V_n - V)$, the solutions are standing waves satisfying $\text{IPR}_2 \propto L^{-1}$. Neglecting finite size effects, a particle initially localised on a single site will spread ballistically from its initial position such that the return probability $|\mathcal{G}_{00}(t)|^2 \propto t^{-1}$. For random V_n , drawn from a nontrivial probability distribution, the wave function at energy E satisfies the discrete Schrödinger equation

$$-t[\psi_{n+1}(E) - \psi_{n-1}(E)] + V_n \psi_n(E) = E \psi_n(E). \quad (1.5.8)$$

Setting $t = 1$ for notational simplicity, the Schrödinger equation may be cast as a recursion relation

$$\begin{pmatrix} \psi_{n+1} \\ \psi_n \end{pmatrix} = \begin{pmatrix} V_n - E & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \psi_n \\ \psi_{n-1} \end{pmatrix}. \quad (1.5.9)$$

The matrix that appears in the above equation is known as the ‘transfer matrix’ and will be denoted $T_n(E)$. The recursion relation (1.5.9) can then be iterated to find the wave function on a given site (given the appropriate boundary conditions, $\psi_1(E)$ and $\psi_0(E)$, at the end of the chain)

$$\begin{pmatrix} \psi_{n+1} \\ \psi_n \end{pmatrix} = T_n T_{n-1} \cdots T_2 T_1 \begin{pmatrix} \psi_1 \\ \psi_0 \end{pmatrix}. \quad (1.5.10)$$

For almost any choice of boundary conditions, the state $\psi_n(E)$ will diverge as $\psi_n \sim e^{n/\xi}$. Equivalently, we may write the following expression for the localisation length

$$\xi^{-1}(E) = \lim_{L \rightarrow \infty} \frac{1}{L} \log \|T_L T_{L-1} \cdots T_2 T_1\| > 0. \quad (1.5.11)$$

Oseledec’s theorem [143, 144] ensures that the limiting matrix $\Gamma = \lim_{L \rightarrow \infty} (Q_L Q_L^\dagger)^{1/(2L)}$ exists, where $Q_L = \prod_{n=1}^L T_n$. The eigenvalues of Γ correspond to the Lyapunov exponents of Q_L . The

smallest of these exponents corresponds to the *largest* length scale, and hence can be identified as the inverse localisation length*. It is worth stressing here that the localisation length ξ corresponds to the *asymptotic* decay of $\psi_\alpha(x)$. This decay need not apply all the way up to maximal value of the wave function, and the “bulk” of the wave function may, in principle, have drastically different properties from its “tails”. It should also be noted that (1.5.11) essentially corresponds to $\langle \ln |t|^2 \rangle$, the average value of the *logarithm* of the transmission coefficient $|t|^2$ associated with the transmission of probability from one end of the chain to the other. It may be shown that $\ln |t|^2$ is approximately normally distributed† and therefore self-averaging, implying that the distribution of $|t|^2$ is approximately log-normal [147]. Conversely, the fat tails of the log normal distribution give rise to peculiar averaging properties for functions of $|t|^2$; whilst the localisation length is self-averaging, the conductance and resistance discussed later on are *not*.

The above discussion shows that in one dimension, *any* amount of spatially uncorrelated disorder will eventually lead to localisation. In the limit of weak disorder $W \ll t$, perturbation theory‡ gives the Thouless expression [145, 149] for the energy-dependent Lyapunov exponent

$$\gamma(E) = \frac{W^2}{24(4t^2 - E^2)} \quad \text{for} \quad E^2 < 4t^2. \quad (1.5.12)$$

This expression contains two important features: the band centre, $E = 0$, is the least strongly localised ($\xi = \gamma^{-1}$ is maximal), and for small disorder the localisation length scales as $\xi \sim W^{-2}$. Conversely, for very strong disorder, the localisation length can be shown to vanish like $\xi^{-1} \simeq \ln \frac{W}{2} - 1$ [150], although more accurate expressions exist [148].

Diagonal versus off-diagonal disorder

“Off diagonal” disorder corresponds to disorder in the hopping amplitudes, rather than in the on-site energies, and can lead to some important differences with respect to conventional Anderson localisation. We will discuss here some of the important differences that arise in one dimension. The one-dimensional tight binding model with pure off-diagonal disorder is

$$\hat{H} = - \sum_n t_{n,n+1} \left(c_n^\dagger c_{n+1} + c_{n+1}^\dagger c_n \right). \quad (1.5.13)$$

*While the discussion thus far has been restricted to strictly one dimensional systems, which possess just one Lyapunov exponent for each energy, analogous arguments apply to *quasi*-one-dimensional systems, whose transfer matrix has multiple Lyapunov exponents [145].

†Technically, one must account for the fact that the transmittance is upper bounded by unity, which gives rise to a *truncated* log-normal distribution [135, 146].

‡An anomaly in the centre of the band leads to a breakdown of second order perturbation theory. The correct coefficient is consistent with $\gamma(0) = W^2/105t^2$, as observed numerically [145, 148].

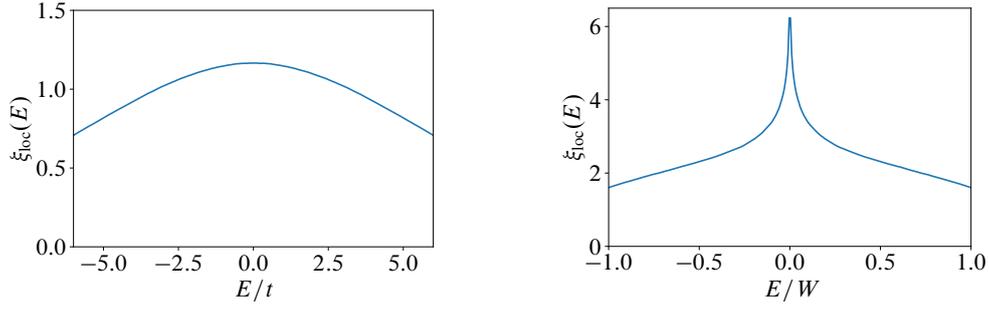


Figure 1.13: A comparison of diagonal (left) versus off-diagonal (right) disorder in a one dimensional Anderson model. The localisation length is always maximal at the centre of the band, $E = 0$. However, in the case of off-diagonal disorder, the localisation length diverges with system size at the centre of the band, $\xi(0) \sim \sqrt{L}$. The curves are calculated using the transfer matrix technique in a system of size $L = 2 \times 10^5$. For the diagonal case, the on-site potentials are drawn from a rectangular distribution between $[-W, W]$ with $W/t = 6$. In the off-diagonal case the hopping parameters are drawn from the same rectangular distribution.

In transfer matrix form, the discrete Schrödinger equation corresponding to (1.5.13) is

$$\begin{pmatrix} \psi_{n+1} \\ \psi_n \end{pmatrix} = \begin{pmatrix} -E/t_{n,n+1} & -t_{n-1,n}/t_{n,n+1} \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \psi_n \\ \psi_{n-1} \end{pmatrix}. \quad (1.5.14)$$

In the centre of the band, $E = 0$, the matrix simplifies substantially, and one finds that

$$T_n T_{n-1} = \begin{pmatrix} -t_{n-1,n}/t_{n,n+1} & 0 \\ 0 & -t_{n-2,n-1}/t_{n-1,n} \end{pmatrix}. \quad (1.5.15)$$

Therefore, $\psi_{2m} = (-1)^m (t_{2m-2,2m-1}/t_{2m-1,2m}) \cdots (t_{0,1}/t_{1,2})$. The first hint that systems with purely off-diagonal disorder can behave in a drastically different manner to the case of conventional, diagonal disorder is to consider the case of binary hopping with random signs: $t_{n,n+1} = e^{i\pi k_n} t$, with k_n a random integer. In this special case, the state in the centre of the band (if it exists) is truly extended with a randomly fluctuating sign. Similarly, (ii) when the chain is dimerised (the Dyson II model), $t_{2n-1,2n} = t_{2n-1,2n-1}$, the state at the centre of the band (if it exists) is extended [151]. Since the system possesses chiral symmetry, states with energies $\pm E$ are paired. This means that only chains with odd L can host a zero mode.

More generally, in the presence of continuously distributed hopping amplitudes, one obtains a divergence of both the density of states, $\rho(E) \sim |E(\ln^3 E)|$ [152], and the localisation length, $\xi(E) \sim |\ln E^2|$ [153, 154], in the vicinity of $E = 0$. Indeed these two quantities are related through the Herbert-Jones-Thouless formulae [155, 156] for one dimensional systems. Although these formulae appear to suggest that there should exist a delocalised state at the centre of the band (indeed, it was believed for some time that this was the case [148]), we must remember that the thermodynamic

limit $L \rightarrow \infty$ has been taken implicitly. In fact, the state in the centre of the band is localised (the transmittance vanishes for sufficiently long systems), but not in the manner described previously. The average that would normally define the localisation length, $\overline{\log |\psi_L/\psi_0|}/L$, vanishes, implying that one must look instead at the fluctuations (which in most cases provide a subleading correction). It can be shown that the fluctuations grow with system size as $\sim \sqrt{L}$, giving a localisation length $\xi(0) \sim \sqrt{L}$ [148, 157]. Consequently, the transmittance behaves as $\langle \ln |t|^2 \rangle \sim -\sqrt{L}$ [147]. In Fig. 1.13, we plot the localisation length of a one dimensional Anderson model in the presence of diagonal and off-diagonal disorder, which illustrate the anomalous behaviour of $\xi(E)$ at the centre of the band.

The scaling hypothesis

In the one dimensional systems we have been considering above, the conductivity is equal to [147, 158]

$$\sigma = \frac{e^2 L}{\pi \hbar} \frac{1}{e^{2L/\xi} - 1}, \quad (1.5.16)$$

where we have temporarily reinstated e and \hbar . When we speak of *the* conductivity σ , we are referring to its geometric average or, equivalently, its typical value. For sufficiently small $L \ll \xi$, the weak localisation regime, Eq. (1.5.16) corresponds to Ohm's law: $\sigma \simeq \sigma_0$ is approximately independent of L and the resistance grows linearly with the length of the system. However, for larger lengths $L \gtrsim \xi$, multiple scattering events become important. In the strongly localised regime, $\xi \ll L$, the wavelike nature of the particles becomes apparent and the resistance grows *exponentially* with the length of the sample*.

The fundamental assumption that underlies the scaling approach is that there is a single parameter that characterises entirely the localisation properties of the system for each dimensionality d . This parameter is the dimensionless DC conductance [149, 159]

$$g = \frac{\hbar}{e^2} \sigma L^{d-2}, \quad (1.5.17)$$

where σL^{d-2} is the conductance per unit cell. Once this parameter has been identified, the β -function, which is assumed to be exclusively a function of g ,

$$\frac{d \ln g}{d \ln L} = \beta(g), \quad (1.5.18)$$

encapsulates the localisation properties. If $\beta(g)$ is positive, then the conductance increases with L to a finite value. Conversely, if $\beta(g)$ is negative, then the conductance “flows” to zero: the

*Here, and in the remainder of this section, we are neglecting any environmental coupling, which will generically lead to a finite coherence time.

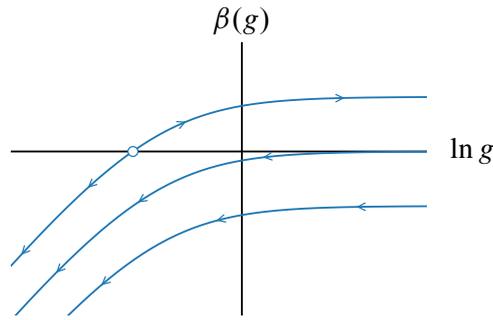


Figure 1.14: Schematic renormalisation group flow diagram for the (dimensionless) conductance, g , $\beta(g) = d \ln g / d \ln L$. For spatial dimensions $d = 1, 2$ the system is always localised; the conductivity flows to zero with system size as $g \sim e^{-L/\xi}$. In $d = 3$, there exists a critical conductance g_c above which the system flows towards an L -independent conductivity.

eigenstates of the system are localised. For sufficiently weak disorder, it is reasonable to assume that the system is essentially metallic, i.e., Ohm's law holds and, consequently, $g \sim L^{d-2}$, giving $\beta(g) \simeq d - 2$. On the other hand, when disorder is very strong, the conductivity should decay exponentially with L , giving $\beta(g) \simeq \ln g$. The qualitative behaviour of the β -function is shown in 1.14, where we have plotted a monotonic function that smoothly interpolates between the two limits. For $d \leq 2$, $\beta(g) < 0$ and the system always flows towards localisation. For $d = 3$, there exists a critical value of the conductance, above which the system is metallic. As we have discussed previously, a finite conductivity does not preclude the existence of localised states, one merely requires that there exists a finite density of extended states for the system to exhibit transport (see Fig. 1.11).

1.5.2 Many-body localisation

The question of whether the phenomenon of Anderson (i.e., single particle) localisation survives the addition of interactions between the constituent degrees of freedom is an old question, prompted by the seminal work of Anderson [126] and coauthors [160]. As discussed in the previous section, one of the defining features of an Anderson insulator is the absence of transport; particles remain exponentially localised in the vicinity of their initial positions. Put alternatively, the state of the system at asymptotically long times is not featureless, but bears some semblance of the system's initial state. This behaviour is at odds with the notion of thermalisation introduced in Sec. 1.3.2, whereby the asymptotic behaviour of a small subsystem is described by a thermal density matrix that depends only on a small number of global conserved quantities [25]. Only recently has the extension to interacting systems been answered definitively: In one dimensional systems subject to strong, quenched randomness, the existence of *many-body* localisation (MBL) can be proven under minimal assumptions [161]. Just like the Anderson insulator, the MBL phase is characterised by the absence of transport [162–166] and evades thermalisation [167–169], retaining some memory of its initial conditions, which can be probed using local observables. Many body localisation is however much more generic than Anderson localisation; the MBL phase can truly be considered a phase of matter, surviving the addition of generic (albeit sufficiently small) perturbations. In the following, we will describe some of the salient features of the MBL phase, starting with the *emergent* conserved quantities known as l-bits that allow MBL systems to evade thermalisation.

Local integrals of motion: l-bits

The canonical model of many body localisation is the spin-1/2 XXZ model in the presence of uncorrelated disorder in the magnetic field in the z direction:

$$\hat{H}_{\text{XXZ}} = J \sum_{i=1}^{L-1} (\hat{S}_i^x \hat{S}_{i+1}^x + \hat{S}_i^y \hat{S}_{i+1}^y) + V \sum_{i=1}^{L-1} \hat{S}_i^z \hat{S}_{i+1}^z + \sum_{i=1}^L h_i \hat{S}_i^z, \quad (1.5.19)$$

where $\hat{S}_i^\alpha = \frac{1}{2} \hat{\sigma}_i^\alpha$ are spin-1/2 degrees of freedom, and h_i is the magnetic field on site i , which is conventionally assumed to be uniformly distributed on the interval $[-W, W]$, independently of its neighbours. Upon performing a standard Jordan–Wigner transformation, the spin Hamiltonian (1.5.19) is mapped onto a number-conserving fermionic Hamiltonian with nearest neighbour interactions

$$\hat{H}_{\text{XXZ}} = \frac{J}{2} \sum_{i=1}^{L-1} (\hat{c}_i^\dagger \hat{c}_{i+1} + \text{H.c.}) + V \sum_{i=1}^{L-1} (\hat{n}_i - \frac{1}{2})(\hat{n}_{i+1} - \frac{1}{2}) + \sum_{i=1}^L h_i (\hat{n}_i - \frac{1}{2}), \quad (1.5.20)$$

with $\hat{n}_i = \hat{c}_i^\dagger \hat{c}_i$ the number of fermions on site i . The anisotropy V maps onto the strength of the nearest neighbour interactions in the fermionic language, while the random magnetic field maps to a random chemical potential. The conservation of fermion number derives from the conservation of the magnetisation in the z direction, $\sum_i \hat{S}_i^z$, in the spin language.

In the special case $V = 0$, the Hamiltonian (1.5.20) (the fermionised XY model) exhibits Anderson localisation of the single-particle wave functions, as explained in Sec. 1.5.1. Alternatively, we can look at the limit $J = 0$, where the Hamiltonian is trivially diagonalised by eigenstates of the \hat{n}_i operators, i.e., local tensor product states $|\{\sigma_i^z\}\rangle$ in the spin language. As J is increased from zero, different $|\{\sigma_i^z\}\rangle$ states hybridise to form eigenstates of (1.5.19). As long as J remains sufficiently small, so that the system remains in the fully many-body localised regime, it is possible to write the Hamiltonian in the form [27, 170]

$$\hat{H}_{\text{XXZ}} = \sum_i h_i \hat{\tau}_i^z + \sum_{i,j} J_{ij} \hat{\tau}_i^z \hat{\tau}_j^z + \sum_{i,j,k} J_{ijk} \hat{\tau}_i^z \hat{\tau}_j^z \hat{\tau}_k^z + \dots \quad (1.5.21)$$

where we have ignored a constant energy shift. The L spin-1/2 operators $\hat{\tau}_i^z$ commute with one another, $[\hat{\tau}_i^z, \hat{\tau}_j^z] = 0$, are conserved quantities, $[\hat{\tau}_i^z, \hat{H}] = 0$, and, crucially, are related to the physical spins $\hat{\sigma}_i^\alpha$ by a quasi-local transformation [170–172]. The operators $\hat{\tau}_i^z$ are known as l-bits, and the Hamiltonian (1.5.21) is able to describe intuitively much of the phenomenology that defines the MBL phase. Formally, the expansion in (1.5.21) is *always* possible (even in the thermalising phase, for example) [39], but the utility of the l-bit representation stems from the quasi-local nature of the transformation between the physical spins and the l-bits. In the thermal phase, the conserved quantities are highly nonlocal, and any physical observable will not have appreciable overlap with them [see however the discussion about quantum many body scars at the end of Sec. 1.3.2, which provide a (fragile) counterexample]. Specifically, the l-bits can be written as [173]

$$\hat{\tau}_i^z = \hat{\sigma}_i^z + \sum_{j,k} \sum_{\alpha,\beta \in \{x,y,z\}} f_{i;jk}^{\alpha\beta} \hat{\sigma}_j^\alpha \hat{\sigma}_k^\beta + \dots, \quad (1.5.22)$$

where the coefficients that define the transformation decay exponentially with the distance from site i , i.e., $f_{i;jk}^{\alpha\beta} \sim e^{-\max(|i-j|, |i-k|)/\xi}$. The length scale ξ may be thought of as a generalisation of the single-particle localisation length [170]. The interactions between the l-bits in (1.5.21) (which are generically absent for Anderson localised systems, which by definition map to free fermions) also decay exponentially in space [170]

$$J_{ij} \sim J_0 e^{-|i-j|/\tilde{\xi}}, \quad (1.5.23)$$

with similar relations holding for multi-l-bit interactions. Note that the dephasing length scale $\tilde{\xi}$ and the length scale ξ controlling the decay of coefficients in the transformation between physical-

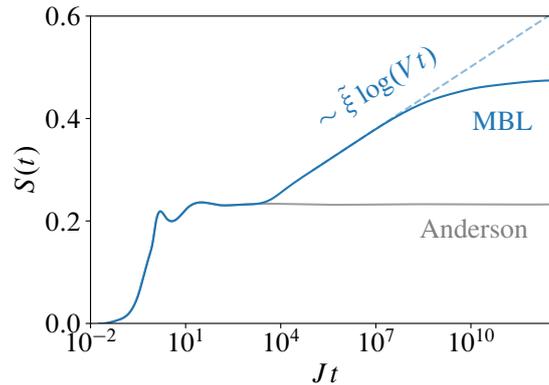


Figure 1.15: Schematic depiction of the growth of the von Neumann entanglement entropy after a quench from a low-entanglement state in the XXZ model (1.5.19). In both cases, there is some initial dynamics that rapidly entangles degrees of freedom separated by $\sim \xi$. In the Anderson localised system, there is no further entanglement growth and the asymptotic state of the system is area law entangled, $S(\infty) \propto \text{const}$. Conversely, in the MBL case, after a certain time (set by the strength of the interactions V) there is an onset of logarithmic entanglement growth. Eventually, the logarithmic growth is truncated by the finite size of the system, and the asymptotic state of the system is volume law entangled, $S(\infty) \propto L$.

and l-bits in (1.5.22) are not necessarily equivalent. Indeed, the two length scales are expected to become increasingly disparate as the MBL transition is approached [174].

Entanglement growth

One characteristic that is able to distinguish between an Anderson insulator and an MBL system is the growth of entanglement following a quench [170, 171, 175–178]. In order to define the entanglement, as we discussed in Sec. 1.2, we must first partition our system into (at least) two distinct subsystems, A and B . In order to observe the slow dephasing generated by nonzero interactions, one must partition the system in such a way that the ‘bulk’ of A and the ‘bulk’ of B are spatially separated*. A particularly convenient choice is to split the system into two equal (and *contiguous*) halves. The system is then initialised in a low-entanglement (e.g., product) state that is not an eigenstate of the Hamiltonian. The subsequent time evolution leads to the generation of entanglement between the two subsystems.

In both Anderson localised and MBL systems, there is some initial dynamics that entangles degrees of freedom separated by $\sim \xi$, the single particle localisation length or the support of the l-bits, respectively. In an Anderson localised system, there is no further dynamics of the entanglement, and one observes a plateau that scales with the length of the boundary separating the two subsystems, $L^{d-1} = \text{const}$ in one dimension. However, in an MBL system, the presence of interactions J_{ij}

*More precisely, the spatial separation of the two subsystems A and B must be large with respect to the length scale ξ . Similarly, by ‘bulk’ we mean the lattice sites that are situated at distances much greater than ξ from the cut(s) between A and B .

between the l-bits in Eq. (1.5.22) leads to entanglement between spins that are separated by distances greater than $\tilde{\xi}$. Specifically, the interaction energy scale $J_{ij} \propto J_0 e^{-|i-j|/\tilde{\xi}}$ between two distant spins at sites i and j , with $|i-j| = \ell$, sets a timescale $t(\ell) = J_0^{-1} e^{\ell/\tilde{\xi}}$ over which entanglement is able to build up between the two spins [178]. Therefore, there exists a logarithmic light cone, $\ell(t) \sim \tilde{\xi} \log(J_0 t)$, inside of which the physical spins are able to become appreciably entangled. Consequently, one observes a logarithmic-in-time growth of the entanglement entropy $S(t) \sim \tilde{\xi} \log(J_0 t)$ [170, 171, 175–178]. Eventually, this logarithmic growth is cut off by system size L , implying volume law scaling at long times. The asymptotic value of the entanglement entropy in the plateau is however typically significantly smaller than would be expected in a typical mid-spectrum state of an ergodic system [173, 178] (i.e., the Page value [179] for random pure states). The difference in behaviour between Anderson and many-body localised systems is shown schematically in Fig. 1.15.

1.5.3 Disorder-free localisation

One of the hallmarks of (many-body) localised phases of matter is the absence of transport. As we have noted previously in Sec. 1.5.1, localisation (both single particle and many-body) is most prevalent in low-dimensional systems, and occurs in the presence of sufficiently strong quenched disorder provided by external potentials or fields. Whether the presence of quenched disorder is not just sufficient but a *necessary* prerequisite for the absence of thermalisation has been a long-standing question. An alternative mechanism by which an isolated, *translationally invariant* quantum system fails to reach local equilibrium starting from generic initial conditions has been dubbed disorder free localisation [180, 181]. In such systems, there exist an extensive number of local conserved quantities (or, more generally, approximately conserved quantities). These conserved quantities, at least within a typical sector, act as an emergent localising potential for the remaining degrees of freedom.

Historically, one of the first models where it was argued that localisation may be induced by interactions alone was put forward by Kagan and Maksimov [182] to describe ^3He defects in solid ^4He . More recently, as part of the broader drive to better understand the thermalisation of closed quantum systems and its absence in many-body localised systems, a number of works have investigated whether localisation can occur in translationally invariant systems [183–190]. A common theme amongst these proposals is the existence of two distinct species of particle, with one much “heavier” than the other* (named heavy-light mixtures). Their larger effective mass endows the heavier particles with slow dynamics, and hence they may intuitively be treated within

*In tight binding models, the larger effective mass of the heavy species of particle manifests as a much smaller tunneling amplitude than the light species. A limiting case is the Falicov–Kimball model, in which one species has no dynamics [191].

a Born–Oppenheimer-like approximation. If the two species interact with one another, the heavy particles can act as a quasi-static disordered background, which localises the light species. A similar phenomenon occurs in (unconstrained [192]) lattice gauge theories [193–199], where static, local conserved charges assume the role of the heavy particles.

It appears that the absence of thermalisation in such translationally invariant systems only persists asymptotically in fine tuned limits where the heavy particles remain precisely static. In the presence of generic perturbations that break the conservation laws (making the heavy particles, in principle, mobile), one expects eventual thermalisation [186, 200]. However, if the perturbations are sufficiently small, memory of the system’s initial conditions can persist over a wide range of intermediate time scales, a phenomenon dubbed *quasi*-MBL [190].

2 | Entanglement negativity in the toric code

2.1 Motivation

With the modern surge of interest in harvesting the capabilities of quantum mechanical systems to develop new technologies, it has become ever so important to quantify and characterise quantum correlations in physical systems. While this remains a tall order in real systems, substantial progress has been made in recent years at the theoretical level. Several measures of quantum correlations (including, but not limited to, “entanglement”) have been proposed and studied (see, e.g., Refs. [8, 201, 202] for reviews). One of the most successful is the von Neumann entanglement entropy. Similarly to other measures, the entanglement entropy works remarkably well when a system is prepared in a pure state—namely, when the density matrix is a projector onto a single quantum mechanical state. However, it becomes less descriptive once we deal with mixed states, as is the case for systems at finite temperature. Under these circumstances, the von Neumann entropy becomes dominated by classical correlations [203] and it is difficult to distil quantum from classical contributions.

A promising way to address this issue has come of late from a newly introduced measure of entanglement, the so-called logarithmic *entanglement negativity* \mathcal{E} [204, 205] and variants thereof [206, 207]. The negativity uses Peres’ positive partial transpose criterion [208] to detect the separability of a system’s density matrix, and as such it is a proper measure of purely quantum correlations—although it is in general only a necessary but not sufficient condition for separability [206, 209].

Unfortunately, calculating \mathcal{E} in many-body quantum systems is notoriously difficult and analytic results are few. At zero temperature, the negativity has been evaluated for a limited number of sufficiently simple lattice models in one [210–212] and two [213, 214] dimensions, suggesting universality at quantum critical points and unveiling leading area law behaviour respectively. That the negativity is fully universal and scale invariant at quantum critical points has subsequently

been proven rigorously using conformal field theory (CFT) [215, 216]. When temperature is raised from zero, one expects that thermal mixing of eigenstates (“thermal fluctuations”) will lead to a reduction in the entanglement. This intuition has been verified for a large number of *two*-spin systems*, initially using the concurrence [217–228] (for a review, see Ref. [201]) and more recently using the negativity [229–231]. Generically, there exists a well-defined temperature T_c , dubbed the “sudden death temperature,” above which the negativity vanishes identically. Beyond two-spin systems, in particular considering the entanglement between two blocks of spins, far less is known. At quantum critical points, CFT results have been extended into the regime of finite temperature [232, 233], while on the numerical side a linked cluster expansion has been employed, alongside exact diagonalisation, to study the negativity at finite temperature in one- and two-dimensional bipartite spin systems [203]. Even though the negativity does not capture *all* quantum correlations, sudden death is nevertheless an intriguing phenomenon that requires further investigation. Further results in $d > 1$ systems, and in particular exact expressions for the negativity at finite temperature, would be highly beneficial to gain a better understanding about the fate of quantum correlations and the origin of the sudden death behaviour.

The goal of this chapter is to present the first exact calculation of the negativity in a 2D lattice system in thermal equilibrium at finite temperature. We choose to work with the toric code model [46] whose exact solubility has allowed the negativity at $T = 0$ [213, 214] and the von Neumann entropy at finite temperature [234] to be calculated exactly. We consider the entanglement between two subsystems that share a boundary of finite length, which allows us to discern the fate of the zero-temperature area law, and how the sudden death temperature T_c depends on boundary length.

2.2 Mixed state separability

2.2.1 PPT criterion

In general, there does not exist a readily computable measure of entanglement whose vanishing is both a necessary *and* sufficient condition for separability as defined in (1.2.10). Instead, we must rely on approximate tests such as the Peres positive partial transpose (PPT) criterion [208], which is only a necessary condition for separability (equivalently, it is only a sufficient condition for entanglement).

Suppose that we would like to test the separability between two disjoint subsystems A_1 , A_2 of a composite system S . We allow for the possibility that A_1 and A_2 may not necessarily cover

*In fact, these studies show that entanglement is *not* always a monotonic function of temperature. Particularly in the presence of a large magnetic field, it is possible for entanglement, as measured by the concurrence, to initially *increase* with increasing temperature before eventually vanishing at T_c , see e.g. Refs. [217, 218].

the whole of S by introducing the complementary subsystem B (which could be the empty set), such that $S = A_1 \cup A_2 \cup B$. Measurements on subsystem $A = A_1 \cup A_2$ are described by the – in general mixed – reduced density matrix $\hat{\rho}_A = \text{Tr}_B \hat{\rho}$. Given an (arbitrary) orthonormal basis for the tensor product Hilbert space of subsystem A , $|\psi_i \phi_j\rangle \equiv |\psi_i\rangle \otimes |\phi_j\rangle \in \mathcal{H}_{A_1} \otimes \mathcal{H}_{A_2}$, we define the operation of partial transposition over subsystem A_2 in terms of matrix elements as [205, 208]

$$\langle \psi_i \phi_j | \hat{\rho}_A^{T_2} | \psi_k \phi_l \rangle = \langle \psi_i \phi_l | \hat{\rho}_A | \psi_k \phi_j \rangle . \quad (2.2.1)$$

Given a separable density matrix $\hat{\rho}_A$, its partial transpose is given by

$$\hat{\rho}_A^{T_2} = \sum_i p_i \hat{\rho}_i^{(1)} \otimes \left(\hat{\rho}_i^{(2)} \right)^T , \quad (2.2.2)$$

where T refers to full transposition. Since $(\hat{\rho}_i^{(2)})^T$ is also a valid density matrix, it follows that $\hat{\rho}_A^{T_2}$ is also a separable state, which therefore implies that all of its eigenvalues are nonnegative. Hence, a state whose partial transpose $\hat{\rho}_A^{T_2}$ has *negative* eigenvalues must be entangled.

2.2.2 The entanglement negativity

Given a tripartite system $S = A_1 \cup A_2 \cup B$, the logarithmic negativity is defined in terms of the reduced state $\hat{\rho}_A = \text{Tr}_B \hat{\rho}$ as

$$\mathcal{E} = \ln \left\| \hat{\rho}_A^{T_2} \right\|_1 , \quad (2.2.3)$$

where T_2 denotes partial transposition over the A_2 subsystem and $\|\cdot\|_1 = \text{Tr}|\cdot|$ is the trace norm [205]. The logarithmic negativity quantifies the entanglement between the subsystems A_1 and A_2 . One can verify that it is symmetric in $A_1 \leftrightarrow A_2$ as a good measure of entanglement must be, so that we can equally compute the transpose over the A_1 subsystem: $\|\hat{\rho}_A^{T_1}\|_1$.

Supposing that the partially transposed density $\hat{\rho}_A^{T_2}$ has eigenvalues λ_i , the trace norm is equivalent to $\|\hat{\rho}_A^{T_1}\|_1 = \sum_i |\lambda_i|$, which may alternatively be written as

$$\|\hat{\rho}_A^{T_1}\|_1 = \sum_{\lambda_i > 0} \lambda_i - \sum_{\lambda_i < 0} \lambda_i = 1 + 2 \sum_{\lambda_i < 0} |\lambda_i| . \quad (2.2.4)$$

It is in this way that the logarithmic negativity quantifies the number of negative eigenvalues λ_i . If all eigenvalues are positive then $\|\hat{\rho}_A^{T_2}\|_1 = 1$, and consequently $\mathcal{E} = 0$, while the existence of *any* negative eigenvalues implies a nonzero negativity $\mathcal{E} > 0$. Hence, the logarithmic negativity quantifies PPT entanglement.

Since the PPT criterion is only a sufficient condition for the presence of entanglement, there exist states which are entangled but nevertheless give a vanishing negativity. However, the negativity

has operational meaning in that it represents an upper bound on the *distillable* entanglement present in a given state [8, 235]. As a result, any entanglement may be present in a state with vanishing negativity $\mathcal{E} = 0$ cannot be distilled, and is therefore referred to as *bound* [235].

2.2.3 The replica method

In order to evaluate the negativity analytically, we employ the replica method, introduced by the authors of Ref. [215] in the context of conformally invariant field theories. The method has since been used to calculate the negativity in a wide variety of models [213, 215, 216, 233, 236, 237, 237]. If we denote the eigenvalues of the partially transposed density matrix $\hat{\rho}_A^{T_2}$ by λ_i , then

$$\mathrm{Tr}(\hat{\rho}_A^{T_2})^{n_e} = \sum_{\lambda_i \geq 0} |\lambda_i|^{n_e} + \sum_{\lambda_i < 0} |\lambda_i|^{n_e}, \quad (2.2.5)$$

$$\mathrm{Tr}(\hat{\rho}_A^{T_2})^{n_o} = \sum_{\lambda_i \geq 0} |\lambda_i|^{n_o} - \sum_{\lambda_i < 0} |\lambda_i|^{n_o}, \quad (2.2.6)$$

where n_e (n_o) is an even (odd) positive integer. In order to obtain the sum of the absolute values of the λ_i , i.e., the trace norm, the replica method consists of following the *even* series and taking the analytic continuation $n_e \rightarrow 1$

$$\mathcal{E} = \lim_{n_e \rightarrow 1} \ln \mathrm{Tr}(\hat{\rho}_A^{T_2})^{n_e}. \quad (2.2.7)$$

If one were to follow instead the odd n series, one would simply obtain the sum of the eigenvalues $\sum_i \lambda_i = \mathrm{Tr} \hat{\rho}_A^{T_2} = \mathrm{Tr} \hat{\rho}_A = 1$. This property can often serve as a useful check of expressions for $\mathrm{Tr}(\hat{\rho}_A^{T_2})^n$. The second equality follows because partial transposition leaves the diagonal elements of the operator (in real space) unchanged.

2.3 Calculations

We consider a system of $2N$ spins in thermodynamic equilibrium with a thermal reservoir at temperature $T = 1/\beta$ (the Boltzmann constant $k_B = 1$ throughout). The system is therefore described by the canonical density matrix $\hat{\rho} = e^{-\beta \hat{H}}/Z$, where \hat{H} is the toric code Hamiltonian introduced in Sec. 1.4.1

$$\hat{H}_{\mathrm{TC}} = -J_e \sum_s \hat{A}_s - J_m \sum_p \hat{B}_p \equiv -J_e \hat{S} - J_m \hat{P}, \quad (2.3.1)$$

and $Z = \mathrm{Tr} e^{-\beta \hat{H}}$ is the partition function (to ensure proper normalisation, $\mathrm{Tr} \hat{\rho} = 1$). The density matrix can be written in terms of its matrix elements with respect to complete bases $\{|\alpha\rangle\}$ and $\{|\gamma\rangle\}$

as

$$\hat{\rho} = \frac{1}{Z} \sum_{\alpha, \gamma} \langle \alpha | e^{-\beta \hat{H}} | \gamma \rangle | \alpha \rangle \langle \gamma | . \quad (2.3.2)$$

Following the work in Refs. [213, 234, 238–240], we always choose to work with the tensor product basis of eigenstates of the operator $\otimes_i \hat{\sigma}_i^z$. The operators $\hat{B}_p \equiv \prod_{i \in p} \hat{\sigma}_i^z$ are diagonal in this basis: $\hat{P} | \alpha \rangle = M(\alpha) | \alpha \rangle$, where one can interpret the eigenvalue $M(\alpha)$ of $\hat{P} = \sum_p \hat{B}_p$ as a “plaquette magnetisation” equal to the sum of all local magnetisations $M_p \equiv \langle \alpha | \hat{B}_p | \alpha \rangle$. We then introduce the group \mathcal{G} generated by products of star operators \hat{A}_s , introduced in Sec. 1.4.1 in the construction of the toric code’s ground state(s). Only states that differ by the action of a group element $g \in \mathcal{G}$ give nonvanishing matrix elements in (2.3.2), which allows us to write $| \gamma \rangle = g | \alpha \rangle$. Hence [213, 234, 238–240],

$$\hat{\rho} = \frac{1}{Z} \sum_{\alpha} \sum_{g \in \mathcal{G}} e^{\beta J_m M(\alpha)} \langle \alpha | e^{\beta \hat{S}} g | \alpha \rangle | \alpha \rangle \langle \alpha | g . \quad (2.3.3)$$

Following similar considerations, the partition function Z can be written as

$$Z = (\cosh \beta J_e)^N [1 + (\tanh \beta J_e)^N] \cdot \sum_{\alpha} e^{\beta J_m M(\alpha)} \quad (2.3.4)$$

$$\equiv Z_e \cdot Z_m . \quad (2.3.5)$$

Notice that the system as a whole is symmetric upon exchanging star and plaquette operators, and $J_e \leftrightarrow J_m$. Therefore, it would be completely equivalent to choose the tensor product basis of the operator $\otimes_i \hat{\sigma}_i^x$, and so we must have $Z_e(\beta J) \propto Z_m(\beta J)$.

To evaluate the matrix elements in (2.3.3), it is helpful to expand the exponential in terms of its constituent star operators. Remembering that $\hat{S} = \sum_s \hat{A}_s$ and that $\hat{A}_s^2 = \mathbb{1}$,

$$e^{\beta J_e \hat{S}} = \prod_s \left[\cosh(\beta J_e) + \sinh(\beta J_e) \hat{A}_s \right] , \quad (2.3.6)$$

and we thus obtain

$$\frac{1}{Z_e} e^{\beta J_e \hat{S}} = \sum_{\tilde{g} \in \mathcal{G}} \eta_T(\tilde{g}) \tilde{g} . \quad (2.3.7)$$

Here we introduced for convenience of notation the weighting factor for each group element

$$\eta_T(g) = \frac{e^{-K_e n(g)} + e^{-K_e [N - n(g)]}}{1 + e^{-K_e N}} , \quad (2.3.8)$$

where $K_e \equiv -\ln \tanh \beta J_e$ and $n(g)$ is the number of star operators \hat{A}_s that appear in the decomposition of the group element g . Note that $\eta_T(g)$ is invariant under $n(g) \rightarrow N - n(g)$, and

therefore the ambiguity in the definition of the group elements of \mathcal{G} modulo the identity is immaterial [213, 234, 238–240]. Making use of (2.3.7), we arrive at the following compact expression for the density matrix of the composite system

$$\hat{\rho} = \frac{1}{Z_m} \sum_{\alpha} \sum_{g \in \mathcal{G}} e^{\beta J_m M(\alpha)} \eta_T(g) |\alpha\rangle \langle \alpha| g. \quad (2.3.9)$$

If we are interested only in a subsystem A of the total system S , then we should trace out the complementary subsystem B to form the reduced density matrix $\hat{\rho}_A = \text{Tr}_B \hat{\rho}$. To achieve this, we decompose $|\alpha\rangle = |\alpha_A\rangle \otimes |\alpha_B\rangle$ and $g = g_A \otimes g_B$,

$$\hat{\rho}_A = \frac{1}{Z_m} \sum_{\alpha} \sum_{g \in \mathcal{G}} e^{\beta J_m M(\alpha)} \eta_T(g) \langle \alpha_B | g_B | \alpha_B \rangle |\alpha_A\rangle \langle \alpha_A| g_A. \quad (2.3.10)$$

Now, the matrix element equals unity if g acts trivially on the B subsystem, and zero otherwise. The resulting constraint on g is implemented by restricting the summation to group elements belonging to the subgroup $\{g \in \mathcal{G} \mid g_B = \mathbb{I}_B\} \equiv \mathcal{G}_A \subset \mathcal{G}$.

Finally, we recall that A is further partitioned into two subsystems ($A = A_1 \cup A_2$), and that we want to take the partial transpose over one of them (say, A_2). This is effected by splitting up the states $|\alpha_A\rangle = |\alpha_{A_1}\rangle \otimes |\alpha_{A_2}\rangle$, and similarly for the group elements $g_A = g_{A_1} \otimes g_{A_2}$

$$\hat{\rho}_A^{T_2} = \frac{1}{Z_m} \sum_{\alpha} \sum_{g \in \mathcal{G}_A} e^{\beta J_m M(\alpha)} \eta_T(g) (|\alpha_{A_1}\rangle \langle \alpha_{A_1}| g_{A_1}) \otimes (g_{A_2} |\alpha_{A_2}\rangle \langle \alpha_{A_2}|). \quad (2.3.11)$$

2.3.1 Star plaquette pair

Up to this point, the manipulations have been completely general. Let us begin by considering the smallest possible subsystem, the star plaquette pair (SPP), depicted in Fig. 2.1. This allows us to gain some intuition about the behaviour of the negativity at nonzero temperatures without any of the complications introduced by considering more elaborate partition schemes. The fact that it is not possible to consider smaller subsystems is a direct consequence of the four-body interactions in the system's Hamiltonian (1.4.1). The SPP system offers a number of technical advantages. Firstly, the negativity is both a necessary *and* sufficient condition for separability*. This allows us to study the entanglement properties of the system exactly. Also, the SPP is symmetric under interchange of stars with plaquettes, implying that the negativity must be symmetric under interchange of $J_e \leftrightarrow J_m$. Finally, one can diagonalise the partially transposed density matrix explicitly without having to resort to the replica trick.

*Details of this result are given in our paper Ref. [241]. The result holds because the negativity is both a necessary and sufficient condition for entanglement when the Hilbert space $\mathcal{H} = \mathbb{C}^2 \otimes \mathbb{C}^2$ [242].

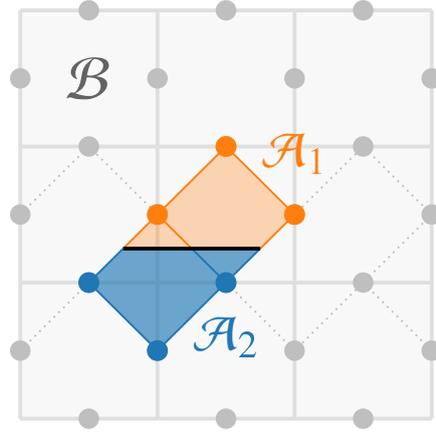


Figure 2.1: Illustration of the star plaquette pair (SPP) subsystem, the smallest subsystem that exhibits a nonvanishing negativity.

For the SPP subsystem, Eq. (2.3.11) simplifies substantially since there exists only one star operator that acts simultaneously on A_1 and A_2 . Therefore, the group \mathcal{G}_A contains only one nontrivial element: the boundary star operator \hat{A}_s . The spins belonging to subsystem B in (2.3.11) can then be “integrated out” by introducing the factors $\xi_T(M_\partial) = \sum_{\alpha_B} e^{\beta J_m M_{\bar{\partial}}(\alpha)} / Z_m$, and by splitting the total plaquette magnetisation into a boundary component $M_\partial \equiv \sum_{p \in \partial} \langle B_p \rangle$, the sum of all boundary plaquettes straddling the A_1 - A_2 boundary, and its complement, so that $M(\alpha) = M_\partial(\alpha_A) + M_{\bar{\partial}}(\alpha)$. Note that the boundary magnetisation depends only on spins contained within subsystem A , while $M_{\bar{\partial}}$ depends on all spins. With the introduction of the ξ_T factors, the partially transposed density matrix becomes

$$\hat{\rho}_A^{T_2} = \sum_{\alpha_A} \sum_{g \in \mathcal{G}_A} e^{\beta J_m M_\partial(\alpha_A)} \xi_T(M_\partial) \eta_T(g) (|\alpha_{A_1}\rangle \langle \alpha_{A_1}| g_{A_1}) \otimes (g_{A_2} |\alpha_{A_2}\rangle \langle \alpha_{A_2}|). \quad (2.3.12)$$

Decomposing the one nontrivial group element into components that act separately on subsystems A_1 and A_2 , $\hat{A}_s = \hat{A}_s^{(1)} \otimes \hat{A}_s^{(2)}$, Eq. (2.3.12) can be written as

$$\hat{\rho}_A^{T_2} = \sum_{\alpha_A} e^{\beta J_m M_\partial(\alpha_A)} \xi_T(M_\partial) \left(|\alpha_A\rangle \langle \alpha_A| + \eta_e |\alpha_{A_1}\rangle \langle \alpha_{A_1}| \otimes \hat{A}_s^{(2)} |\alpha_{A_2}\rangle \langle \alpha_{A_2}| \right), \quad (2.3.13)$$

where $\eta_e \equiv \eta(\hat{A}_s)$, i.e., $n(g) = 1$ in (2.3.8). Relabelling the states α in the second term, the density matrix may be written more compactly as

$$\hat{\rho}_A^{T_2} = \sum_{\alpha_A} e^{\beta J_m M_\partial(\alpha_A)} \xi_T(M_\partial) |\alpha_A\rangle \langle \alpha_A| + \eta_e e^{-\beta J_m M_\partial(\alpha_A)} \xi_T(-M_\partial) |\alpha_A\rangle \langle \alpha_A| \hat{A}_s. \quad (2.3.14)$$

Each state $|\alpha_A\rangle$ of the subsystem A is coupled only to itself and its counterpart with the spins $i \in s$ flipped. The expression (2.3.14) implies that $\hat{\rho}_A^{T_2}$ can be written as a block diagonal matrix with

respect to the eigenstates of $\otimes_{i \in A} \hat{\sigma}_i^z$, with the 2×2 matrices

$$\rho_{\pm} = \begin{pmatrix} \Xi_m^{\pm} & \eta_e \Xi_m^{\mp} \\ \eta_e \Xi_m^{\mp} & \Xi_m^{\pm} \end{pmatrix} \quad (2.3.15)$$

along the diagonal (note that these are not density matrices). For convenience, we introduced the shorthand Ξ_m^{\pm} to represent $\xi_T(M_{\partial})e^{\beta J_m M_{\partial}}$ evaluated with boundary magnetisation $M_{\partial} = \pm 1$. Since the number of spin configurations α_A with $M_{\partial} = +1$ equals the number with $M_{\partial} = -1$, there exist an equal number of ρ_+ and ρ_- blocks along the diagonal of $\hat{\rho}_A^{T_2}$. Therefore the normalisation condition $\text{Tr} \hat{\rho}_A^{T_2} = 1$ gives $\Xi_m^+ + \Xi_m^- = 2^{1-\mathcal{N}_A}$, where $\mathcal{N}_A \equiv 6$ equals the number of spins $i \in A$.

After calculating the eigenvalues of ρ_{\pm} , invoking the symmetry $\mathcal{E}(J_e, J_m) = \mathcal{E}(J_m, J_e)$ allows us to conjecture that $\Xi_m^+ - \Xi_m^- = 2^{1-\mathcal{N}_A} \eta_m$. This implies that $\Xi_m^{\pm} = (1 \pm \eta_m)/2^{\mathcal{N}_A}$, which leads to the final result for the negativity of the SPP subsystem

$$e^{\mathcal{E}(T)} = 1 + \frac{1}{2} \max(\eta_e + \eta_m + \eta_e \eta_m - 1, 0), \quad (2.3.16)$$

plotted in Fig. 2.4. By construction, this expression is symmetric under interchange of $J_e \leftrightarrow J_m$. This expression can also be derived without resorting to symmetry using the techniques described in Sec. 2.3.2.

The result (2.3.16) allows for the following physical interpretation. When temperature is finite, but small in comparison to J_e , the densities of both species are exponentially suppressed. The bath is unable to provide sufficient thermal energy to excite either species of defect with any significant probability. As a result, the negativity remains essentially undegraded from its zero temperature value, $\mathcal{E}(T) \simeq \mathcal{E}(0)$, over this temperature range. Namely, a density of order one defect per site of the least energetically costly defects is required to start affecting the quantum correlations between A_1 and A_2 . Above the temperature threshold $T \sim J_e$, electric charge defects proliferate in the system and the negativity correspondingly decays like $\mathcal{E} \sim J_e/T$. This decay is cut off at a temperature $T_c \sim J_m$ (up to logarithmic factors discussed below), where the negativity exhibits sudden death. This temperature corresponds to the excitation of an $O(1)$ density of the most energetically costly defects. The precise expression for the sudden death temperature is found by solving $\eta_e + \eta_m + \eta_e \eta_m = 1$, which, in the thermodynamic limit $N \rightarrow \infty$, reduces to

$$n_h(2J_e)n_h(2J_m) = \frac{1}{2}. \quad (2.3.17)$$

We have introduced the density of holes $n_h(2J_X) = 1 - n_F(2J_X)$, i.e., the density of nondefective stars (plaquettes) for $X = e (m)$. At T_c the geometric mean of the hole densities must equal $1/\sqrt{2}$. When the coupling constants are well separated, we can obtain an approximate explicit expression

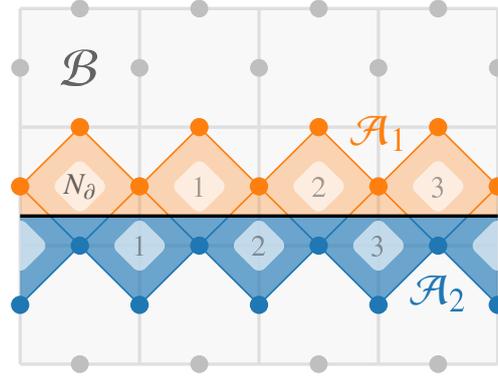


Figure 2.2: Illustration of the noncontractible “strip” partition considered when calculating the dependence of the negativity on the length of the A_1 - A_2 boundary.

for the sudden death temperature

$$T_c = \frac{2J}{W(J)} \sim \frac{2J}{\ln J} \quad \text{for } J \gg 1, \quad (2.3.18)$$

where where $J \equiv J_m/J_e$ and $W(x)$ is the product-log function. Using the duality $J_e \leftrightarrow J_m$, we can also use (2.3.18) to deduce that T_c vanishes logarithmically as $J_e \rightarrow 0$.

2.3.2 Extended boundary

The results for the SPP are simple and intuitive, but do not address the question of how the negativity, and in particular the sudden death temperature T_c , depend on the length of the A_1 - A_2 boundary. We devote this section to answering this question.

We focus on the partition scheme depicted in Fig. 2.2, i.e., a “strip” that spans the torus in one direction. This is the simplest choice of subsystem that allows one to study the effect of increasing the length of the A_1 - A_2 boundary. One can imagine constructing the strip partition by tiling neighbouring SPP subsystems. Choosing a noncontractible subsystem makes the calculation simpler than for a contractible strip by eliminating its endpoints, and therefore the associated boundary effects. In the interests of brevity, we will only present the results for the noncontractible strip, but it can be shown that the results for the contractible strip are both qualitatively and quantitatively similar in the limit of large boundaries. The one substantial difference between considering a noncontractible versus a contractible subsystem is the presence of subleading corrections to the zero-temperature area law in the former case, i.e., $\mathcal{E}(0) \propto N_\partial - 1$, which indicates the presence of topological order [213].

The replica calculation begins by taking the trace of the n th power of (2.3.11):

$$\begin{aligned} \text{Tr}(\hat{\rho}_A^{T_2})^n &= \frac{1}{Z_m^n} \sum_{\alpha_1, \dots, \alpha_n} \sum_{g_1, \dots, g_n \in \mathcal{G}_A} \left(\prod_{\ell=1}^n e^{\beta J_m M(\alpha_\ell)} \eta_T(g_\ell) \right) \\ &\quad \langle \alpha_{1A_1} | g_{1A_1} | \alpha_{2A_1} \rangle \cdots \langle \alpha_{(n-1)A_1} | g_{(n-1)A_1} | \alpha_{nA_1} \rangle \langle \alpha_{nA_1} | g_{nA_1} | \alpha_{1A_1} \rangle \\ &\quad \langle \alpha_{1A_2} | g_{2A_2} | \alpha_{2A_2} \rangle \cdots \langle \alpha_{(n-1)A_2} | g_{nA_2} | \alpha_{nA_2} \rangle \langle \alpha_{nA_2} | g_{1A_2} | \alpha_{1A_2} \rangle . \end{aligned} \quad (2.3.19)$$

It is helpful to notice that the subgroup \mathcal{G}_A can in general be decomposed as $\mathcal{G}_A = \mathcal{G}_{A_1} \mathcal{G}_{A_2} \mathcal{G}_{A_1 A_2}$. The subgroups \mathcal{G}_{A_i} are defined as $\mathcal{G}_{A_i} \equiv \{g \in \mathcal{G}_A \mid g_{\bar{A}_i} = \mathbb{I}_{\bar{A}_i}\}$ with $\bar{A}_1 = A_2$ and vice versa, while $\mathcal{G}_{A_1 A_2} \equiv \mathcal{G}_A / (\mathcal{G}_{A_1} \mathcal{G}_{A_2})$ is the quotient group. Any element $g \in \mathcal{G}_A$ can therefore be uniquely decomposed into the product of three group elements: one that acts exclusively on A_1 , one exclusively on A_2 , and one that acts simultaneously (and exclusively) on A_1 and A_2 . Namely, $g = \bar{g} \bar{\bar{g}} \theta$, with $\bar{g} \in \mathcal{G}_{A_1}$, $\bar{\bar{g}} \in \mathcal{G}_{A_2}$, and $\theta \in \mathcal{G}_{A_1 A_2}$ [213].

Using this decomposition, we make the following relabelling of the states $|\alpha_k\rangle$, for all $k > 1$:

$$|\alpha'_k\rangle = \left(\prod_{\ell=1}^{k-1} \bar{g}_\ell \right) \left(\prod_{\ell=2}^k \bar{\bar{g}}_\ell \right) |\alpha_k\rangle \rightarrow |\alpha_k\rangle . \quad (2.3.20)$$

Note that this redefinition does not affect the plaquette magnetisation, since the group elements $\bar{g} \in \mathcal{G}_{A_1}$ and $\bar{\bar{g}} \in \mathcal{G}_{A_2}$ flip only closed loops of spins on the dual lattice. This removes all dependence on \bar{g}_ℓ and $\bar{\bar{g}}_\ell$ from the matrix elements, and we obtain a more compact expression:

$$\begin{aligned} \text{Tr}(\hat{\rho}_A^{T_2})^n &= \frac{1}{Z_m^n} \sum_{\alpha_1, \dots, \alpha_n} \sum_{\substack{\bar{g}_1, \dots, \bar{g}_n \\ \in \mathcal{G}_{A_1}}} \sum_{\substack{\bar{\bar{g}}_1, \dots, \bar{\bar{g}}_n \\ \in \mathcal{G}_{A_2}}} \sum_{\substack{\theta_1, \dots, \theta_n \\ \in \mathcal{G}_{A_1 A_2}}} \langle 0 | \prod_{\ell=1}^n \bar{g}_\ell \bar{\bar{g}}_\ell | 0 \rangle \\ &\quad \left(\prod_{\ell=1}^n e^{\beta J_m M(\alpha_\ell)} \eta_T(\bar{g}_\ell \bar{\bar{g}}_\ell \theta_\ell) \langle \alpha_{\ell A} | \theta_{\ell A_1} \otimes \theta_{(\ell+1) A_2} | \alpha_{(\ell+1) A} \rangle \right) , \end{aligned} \quad (2.3.21)$$

where $\theta_{n+1} \equiv \theta_1$ and similarly $\alpha_{n+1} \equiv \alpha_1$. Notice that the matrix elements in Eq. (2.3.21) impose constraints on which terms \bar{g}_ℓ and $\bar{\bar{g}}_\ell$ in the summation give a nonvanishing contribution. Furthermore, they implicitly impose the constraint $\prod_\ell \theta_\ell = \mathbb{I}$. This is because, given α_1 (say), it takes only $n-1$ of the n matrix elements to uniquely determine all other α_ℓ ; the final matrix element then evaluates to $\langle \alpha_1 | \prod_\ell \theta_\ell | \alpha_1 \rangle$, which indeed is nonvanishing only if the aforementioned constraint is satisfied.

The strip partition scheme (Fig. 2.2) offers various simplifications. Both A_1 and A_2 span the torus in one direction and share one edge. The subsystem $A \equiv A_1 \cup A_2$ consists only of boundary stars ($\{\hat{A}_s \mid s \in \partial\}$) and boundary plaquettes ($\{\hat{B}_p \mid p \in \partial\}$) which act simultaneously (and

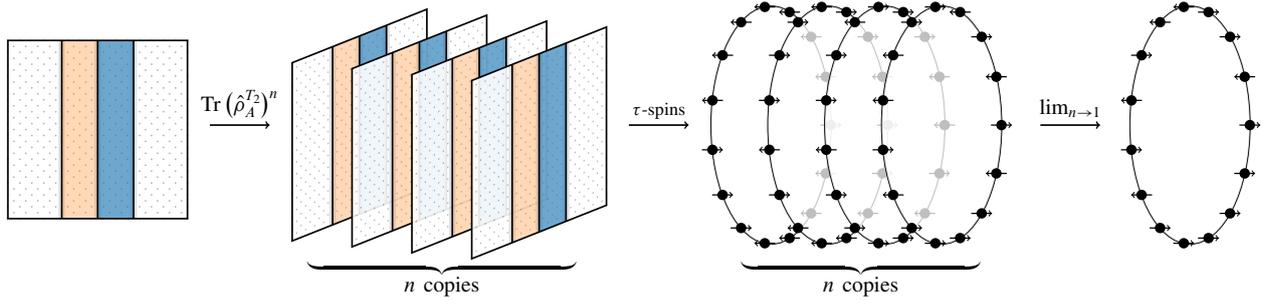


Figure 2.3: A graphical map of the replica approach. Taking the trace of the n th power of $\hat{\rho}_A^{T_2}$ creates n copies of the system. Introducing classical spin variables τ that live on each boundary star maps the system onto n one-dimensional Ising chains. In the limit $n \rightarrow 1$, we arrive at the final result, which is related to the partition function of a single one-dimensional, classical Ising chain.

exclusively) on both the A_1 and A_2 subsystems. We therefore deduce that the groups \mathcal{G}_{A_1} and \mathcal{G}_{A_2} become trivial ($\{\mathbb{1}\}$), removing the summations over group elements \bar{g}_ℓ and $\bar{\bar{g}}_\ell$. In addition, A is symmetric under interchange of stars and plaquettes so, as for the SPP system, the final result must obey the symmetry $\mathcal{E}(J_e, J_m) = \mathcal{E}(J_m, J_e)$.

The magnetisation can be decomposed into a boundary component M_∂ , equal to the sum of the magnetisations of all boundary plaquettes straddling A_1 and A_2 , and its complement: $M(\alpha) = M_\partial(\alpha_A) + M_{\bar{\partial}}(\alpha)$. Similarly to the SPP, tracing over subsystem B is then entirely contained within a factor $\xi_T(\alpha_A) \equiv \frac{1}{Z_m} \sum_{\alpha_B} e^{\beta J_m M_{\bar{\partial}}(\alpha)}$, whose explicit calculation is left until later. After making these simplifications, (2.3.21) becomes

$$\text{Tr}(\hat{\rho}_A^{T_2})^n = \sum_{\alpha_{1A}, \dots, \alpha_{nA}} \sum_{\substack{\theta_1, \dots, \theta_n \\ \in \mathcal{G}_{A_1 A_2}}} \langle 0 | \prod_{\ell=1}^n \theta_\ell | 0 \rangle e^{\beta J_m \sum_{\ell} M_\partial(\alpha_{\ell A})} \prod_{\ell=1}^n \xi_T(\alpha_{\ell A}) \eta_T(\theta_\ell) \langle \alpha_{\ell A} | \theta_{\ell A_1} \otimes \theta_{(\ell+1)A_2} | \alpha_{(\ell+1)A} \rangle. \quad (2.3.22)$$

In order to apply the replica trick, we must make all dependence of the expression (2.3.22) on n explicit, which will allow us to take the limit $n \rightarrow 1$. This approach is depicted schematically in Fig. 2.3. To this end, we introduce the classical variables $\sigma_s^{(\ell)}$ that take the value 0 (1) if the star \hat{A}_s is present (not present) in the decomposition of the boundary group element θ_ℓ . Therefore, by definition, we have the relationship

$$\theta_\ell = \prod_s \hat{A}_s^{\sigma_s^{(\ell)}}. \quad (2.3.23)$$

This implies that the boundary matrix element in (2.3.22) can be written as

$$\langle 0 | \prod_{\ell=1}^n \theta_{\ell} | 0 \rangle = \langle 0 | \prod_{s \in \partial} \hat{A}_s^{\sum_{\ell} \sigma_s^{(\ell)}} | 0 \rangle = \prod_{s \in \partial} \delta \left(\sum_{\ell=1}^n \sigma_s^{(\ell)} \bmod 2 \right). \quad (2.3.24)$$

The second equality holds because the matrix element equals unity if and only if an even number of \hat{A}_s act on each star. It is more convenient to work with the following representation of the Kronecker delta

$$\delta(y \bmod 2) = \frac{1}{2} \sum_{\mu=0}^1 e^{i\pi \mu y}. \quad (2.3.25)$$

The benefit of introducing this alternative representation becomes evident when we take the product over boundary stars $s \in \partial$

$$\langle 0 | \prod_{\ell} \theta_{\ell} | 0 \rangle = \frac{1}{2^{N_{\partial}}} \sum_{\{\mu_s\}} e^{i\pi \sum_{\ell} \sum_s \mu_s \sigma_s^{(\ell)}}. \quad (2.3.26)$$

At the expense of introducing a further N_{∂} variables $\{\mu_s\}$, we have been able to write the constraint (2.3.24) in a form that is separable into a product over the different replicas (similar in spirit to a Hubbard–Stratonovich transformation). In order to represent the weighting factor $\eta_T(\theta_{\ell})$ in terms of the new variables, we need an expression for the number of stars present in the decomposition of the boundary group element θ_{ℓ} . This has a particularly simple representation in terms of $\sigma_s^{(\ell)}$ variables: $n(\theta_{\ell}) = \sum_s \sigma_s^{(\ell)}$, which allows us to write $\eta_T(\theta_{\ell})$ explicitly as

$$\eta_T(\theta_{\ell}) = \frac{1}{2 \cosh\left(\frac{NK_e}{2}\right)} \sum_{J_{\ell}=\pm 1} e^{J_{\ell} K_e (N/2 - \sum_s \sigma_s^{(\ell)})}, \quad (2.3.27)$$

where we recall that $K_e = -\ln \tanh \beta J_e$.

To evaluate the ξ_T weighting factor for the magnetic vortex configurations, it is more convenient to work with configurations of the plaquette magnetisations M_p , the eigenvalues of \hat{B}_p in the state α , rather than the physical spin configurations. This is possible because $\xi_T(\alpha_A) = \xi_T(\alpha'_A)$ if all boundary plaquette magnetisations $\{M_p \mid p \in \partial\}$ of α_A and α'_A are the same. The magnetisations can then be treated as Ising spins that live at the centres of the plaquettes p , subject to the constraint $\prod_p M_p = 1$, which follows from the operator relation $\prod_p \hat{B}_p = \mathbb{1}$. Hence, the only effect of the boundary magnetisation configuration on the *bulk* magnetisation is an even-odd effect: if an even (odd) number of boundary plaquette magnetisations are negative, there must be a compensating even (odd) number of negative plaquette magnetisations in the bulk. This is an artefact of the periodic boundary conditions, which are the source of the constraint $\prod_p M_p = 1$.

Summing over bulk magnetisation configurations $\{M_{p \notin \partial}\}$, keeping the boundary configuration

$\{M_{p \in \partial}\}$ fixed, we arrive at

$$\xi_T(M_\partial) = \frac{1}{2^{3N_\partial}} \frac{\cosh(\beta J_m)^{N-N_\partial} + (\prod_{p \in \partial} M_p) \sinh(\beta J_m)^{N-N_\partial}}{\cosh(\beta J_m)^N + \sinh(\beta J_m)^N}. \quad (2.3.28)$$

In accordance with the arguments made earlier, only the sign of $\prod_{p \in \partial} M_p$ enters into the expression for ξ_T . In the thermodynamic limit, the even-odd effect is removed and one obtains the simplified result

$$\xi_T(M_\partial) = \frac{1}{2^{3N_\partial}} \operatorname{sech}(\beta J_m)^{N_\partial}. \quad (2.3.29)$$

Since each boundary star is adjacent to two boundary plaquettes, the parity of the number of negative $M_{p \in \partial}$ (i.e., the sign of $\prod_{p \in \partial} M_p$) is conserved between replicas so that $\prod_\ell \xi_T(M_\partial(\alpha_\ell)) = \xi_T^n(M_\partial(\alpha_1))$. This simplification is only possible for the noncontractible strip partition, since open ends allow for a parity change between adjacent replicas.

The final step in the conversion to classical spin variables is to find the boundary plaquette magnetisation for each replica, $M_\partial = \sum_{p \in \partial} M_p$. Given α_{1A} , say, then all $\alpha_{\ell A}$ for $\ell > 1$ are uniquely determined by (2.3.22). Introducing Ising spins $\tau_s^{(\ell)} = 2\sigma_s^{(\ell)} - 1 \in \{-1, 1\}$, the recursion relation between boundary magnetisations M_s , $s = p \in \partial$, of adjacent replicas is $M_s(\alpha_\ell) = M_s(\alpha_{\ell-1}) \tau_s^{(\ell)} \tau_{s+1}^{(\ell)} \tau_s^{(\ell-1)} \tau_{s+1}^{(\ell-1)}$. Alternatively, iterating this recursion relation down to the magnetisation of the spin configuration α_1 ,

$$M_s(\alpha_\ell) = M_s(\alpha_1) \tau_s^{(1)} \tau_{s+1}^{(1)} \tau_s^{(\ell)} \tau_{s+1}^{(\ell)}. \quad (2.3.30)$$

The notation M_s is interpreted as the magnetisation of the plaquette with the same index as the star s , as depicted in Fig. 2.2. Periodic boundary conditions are imposed on the τ -spins in real space (i.e., $\tau_1 = \tau_{N_\partial+1}$, separately for each replica ℓ). The description in terms of classical spin variables is now complete and we are in a position to substitute back into (2.3.22). Relabelling $\alpha_1 \rightarrow \alpha$, we obtain

$$\begin{aligned} \operatorname{Tr} (\hat{\rho}_A^{T_2})^n &= \frac{1}{2^{N_\partial}} \frac{1}{[2 \cosh(\frac{NK_e}{2})]^n} \sum_{\alpha_A} \xi_T^n(M_\partial(\alpha_A)) \\ &\quad \sum_{\{\mu_s\}_{s=1}^{N_\partial}} \sum_{\{J_\ell\}_{\ell=1}^n} \prod_{\ell=1}^n \sum_{\{\sigma_s^{(\ell)}\}_{s=1}^{N_\partial}} e^{-\sum_{s=1}^{N_\partial} [\gamma_s \tau_s^{(\ell)} \tau_{s+1}^{(\ell)} + h_s^{(\ell)} \sigma_s^{(\ell)}] + \frac{1}{2} N J_\ell K_e}, \end{aligned} \quad (2.3.31)$$

where $h_s^{(\ell)} \equiv J_\ell K_e + i\pi\mu_s$ and $\gamma_s \equiv -\beta J_m M_s$. Notice that we have performed a relabelling of the physical spin configurations α_A , for each classical spin configuration $\{\tau_s^{(\ell)}\}$, such that $M_s \tau_s^{(1)} \tau_{s+1}^{(1)} \rightarrow M_s$ [the parity is conserved by this transformation since $\prod_s M_s \tau_s^{(1)} \tau_{s+1}^{(1)} = \prod_s M_s$, thus leaving $\xi_T^n(M_\partial)$ unchanged]. This puts all replicas on an equal footing, making the trace

separable into a product over replicas. The problem has therefore been reduced to evaluating the classical partition function

$$Z(\{\gamma_s\}, \{\mu_s\}, J_\ell K_e) = \sum_{\{\tau_s = \pm 1\}} e^{-\sum_{s=1}^{N_\partial} [\gamma_s \tau_s \tau_{s+1} + \frac{1}{2} h_s^{(\ell)} (\tau_s + 1)]} \quad (2.3.32)$$

of a 1D Ising model that lives on the boundary between the regions A_1 and A_2 , with reduced couplings $\{\gamma_s\}$ in a complex reduced magnetic field $\{h_s\}$. Note that because the magnetic field is complex, the partition function is not necessarily positive. Using the expression

$$\lim_{n \rightarrow 1} \sum_{m=0}^n \binom{n}{m} x^m y^{n-m} = |x + y|, \quad (2.3.33)$$

having followed the *even* n series, one can then evaluate the sum over configurations $\{J_\ell\}$

$$\sum_{\{J_\ell\}_{\ell=1}^n} \prod_{\ell=1}^n e^{\frac{1}{2} N J_\ell K_e} Z(J_\ell K_e) = \left| e^{\frac{1}{2} N K_e} Z(K_e) + e^{-\frac{1}{2} N K_e} Z(-K_e) \right|, \quad (2.3.34)$$

where we have suppressed the dependence of the partition function Z on $\{\gamma_s\}$ and $\{\mu_s\}$ to simplify the notation.

At this stage we take the thermodynamic limit $N \rightarrow \infty$. The partition function $Z(-K_e)$ is suppressed by a factor $e^{-N K_e}$, and the even-odd effect is removed from the weighting factor ξ_T . Absorbing ξ_T into the partition function as an energy shift and trading the sum over spin configurations α_A for one over magnetisations $\{M_s\}$, we obtain

$$\|\hat{\rho}_A^{T_2}\|_1 = \frac{1}{2^{2N_\partial}} \sum_{\{\gamma_s\}_{s=1}^{N_\partial}} \sum_{\{\mu_s\}_{s=1}^{N_\partial}} |Z(\{\gamma_s\}, \{\mu_s\})|, \quad (2.3.35)$$

where we have redefined

$$Z = \sum_{\{\tau_s\}} \exp \left\{ - \sum_{s=1}^{N_\partial} \left[\gamma_s \tau_s \tau_{s+1} + \frac{1}{4} h_s (\tau_s + \tau_{s+1} + 2) + \ln \cosh \beta J_m \right] \right\}. \quad (2.3.36)$$

The sums over configurations $\{\gamma_s\} \equiv \gamma$ and $\{\mu_s\} \equiv \mu$ can be physically interpreted as performing a disorder average over the reduced couplings and the reduced magnetic field, respectively. Therefore, we arrive at the following final expression for the negativity of the strip partition scheme in the thermodynamic limit

$$\mathcal{E}(J_e, J_m) = \ln \langle |Z(\gamma, \mu)| \rangle, \quad (2.3.37)$$

where the angled brackets refer to a disorder average over γ and μ . In the limit $T \rightarrow 0$, this

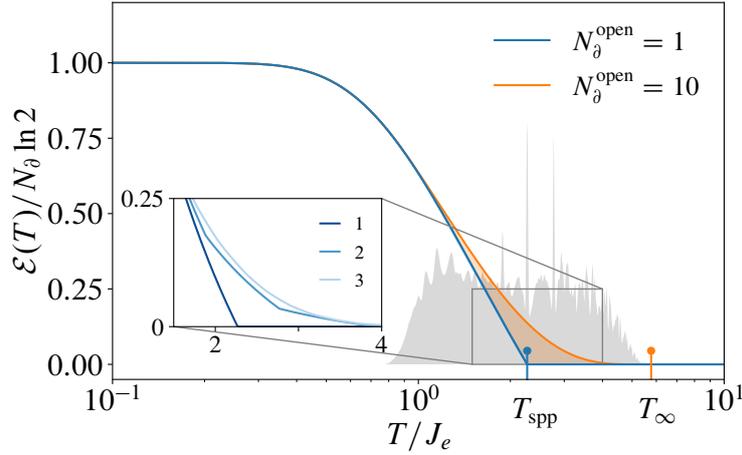


Figure 2.4: Negativity for a contractible strip partition, as discussed in the text, with boundary length N_∂^{open} and $J_e = J_m$. As the boundary between A_1 and A_2 is made longer, \mathcal{E} develops an increasing number of discontinuities (at temperatures T_i) in its first derivative whose effect is to “smooth” the function $\mathcal{E}(T)$ and push T_c to higher temperatures. In the background is plotted the “density of discontinuities” $\sum_i \delta(T - T_i)$ for $N_\partial^{\text{open}} = 7$. The vertical lines indicate the positions of T_c for the SPP and for an infinite boundary. The zoomed inset shows the negativity for $N_\partial^{\text{open}} = 1, 2, 3$, highlighting the rapidly increasing number of discontinuities.

expression reduces to $\mathcal{E}(0) = (N_\partial - 1) \ln 2$, in agreement with the results of Ref. [213, 214]. The result is plotted as a function of temperature in Fig. 2.4. The expression (2.3.34) has been shown to agree precisely with exact diagonalisation results. As temperature is increased from zero, similarly to the SPP, the negativity remains undegraded, $\mathcal{E}(T) \simeq \mathcal{E}(0)$, until the density of thermally excited electric defects becomes of $O(1)$. It then vanishes identically above some temperature T_c , discussed in more depth below. The SPP and extended cases however differ in that the latter exhibits a large number of discontinuities. These discontinuities correspond to the (Lee-Yang [243, 244]) zeros of the complex partition functions in (2.3.35), whose effect is to “smooth out” $\mathcal{E}(T)$, and to push T_c to higher temperatures. The number of discontinuities grows rapidly (exponentially for $N_\partial \gg 1$) with increasing boundary length, e.g., 1, 19, 78, ... for $N_\partial = 2, 3, 4, \dots$ and $J_e = J_m$.

One can show that performing the disorder average *without* taking the absolute value first results in $\langle Z \rangle = 1$. Therefore, if *all* disorder realisations produce a positive partition function, then $\langle |Z| \rangle = \langle Z \rangle$, and the negativity vanishes identically $\mathcal{E} = 0$. This is the case at infinite temperature where $Z(\gamma, \mu) \rightarrow 1$ for all disorder realisations. If at least one disorder realisation becomes negative, then the average satisfies $\langle |Z| \rangle > 1$, and consequently the negativity becomes nonvanishing $\mathcal{E} > 0$. For each disorder realisation, there exists some temperature, which we will refer to as T_* , above which $Z(\gamma, \mu) > 0$. The sudden death temperature T_c corresponds to the maximum of the temperatures T_* over all disorder realisations.

The disorder realisation that maximises T_* corresponds to a special *ordered* configuration of the parameters γ, μ . In particular, we find that $\gamma_s = \beta J_m$ and $\mu_s = 1, \forall s$ give rise to the partition

function with the largest zero. Because the resulting Ising chain is translationally invariant, one can use a transfer matrix approach to calculate the partition function analytically. The partition function is written as a trace over a product of transfer matrices \mathbb{T}_ℓ

$$Z(\gamma, \mu) = \text{Tr} \left[\prod_{\ell=1}^{N_\partial} \mathbb{T}_\ell \right]. \quad (2.3.38)$$

Because the couplings and magnetic field are spatially homogeneous, all transfer matrices are identical, $\mathbb{T}_\ell = \mathbb{T}$, $\forall \ell$, which can be chosen* to assume the symmetric form

$$\mathbb{T} = e^{-h/2} \begin{pmatrix} e^{-\gamma-h/2} & e^\gamma \\ e^\gamma & e^{-\gamma+h/2} \end{pmatrix}. \quad (2.3.39)$$

In order to find the sudden death temperature, we must solve the equation $Z = 0$. This equation has multiple solutions as the partition function oscillates with temperature. The *highest* of these temperatures corresponds to the sudden death temperature T_c , above which the negativity vanishes identically. Substituting $h = K_e + i\pi$, the matrix \mathbb{T} has eigenvalues

$$\lambda_{\pm} = \frac{i}{2} e^{-\gamma} \left[e^{K_e} - 1 \pm i \sqrt{e^{K_e} (4e^{4\gamma} - e^{K_e} - e^{-K_e} - 2)} \right]. \quad (2.3.40)$$

In terms of these eigenvalues, the condition for sudden death becomes

$$\lambda_+^{N_\partial} + \lambda_-^{N_\partial} = 0. \quad (2.3.41)$$

Rearranging this expression,

$$\left(\frac{\lambda_+}{\lambda_-} \right)^{N_\partial} \equiv \left(\frac{1+ix}{1-ix} \right)^{N_\partial} = -1. \quad (2.3.42)$$

This equation admits the following solutions for x , which are parameterised by the odd integer m

$$x = \tan \left(\frac{m\pi}{2N_\partial} \right). \quad (2.3.43)$$

The solutions of this equation that occur at temperatures lower than the sudden death temperature correspond to the locations of further discontinuities in the derivative of $\mathcal{E}(T)$. Using the expressions for the eigenvalues (2.3.40), the explicit form of x in terms of the coupling constants J_e and J_m is

$$x = \sqrt{e^{4\beta(J_e+J_m)} - e^{4\beta J_m} - e^{4\beta J_e}}. \quad (2.3.44)$$

*There is some freedom in how one writes the linear term, $\sum_i h_i \tau_i$, which corresponds to the fictitious magnetic field. The most symmetric choice is to write this term as $\sum_i (h_i \tau_i + h_{i+1} \tau_{i+1})/2$, as we have done in (2.3.36).

The integer m which gives rise to the largest (real) temperature solution corresponds to the smallest positive x . Therefore, the condition for sudden death may be written explicitly as

$$e^{4\beta(J_e+J_m)} - e^{4\beta J_m} - e^{4\beta J_e} = \tan^2\left(\frac{\pi}{2N_\partial}\right). \quad (2.3.45)$$

The expression obeys the symmetry $J_e \leftrightarrow J_m$, as it must.

In a similar manner to the SPP, the sudden death criterion may be written in terms of the defect densities as

$$\frac{n_F(2J_e)n_F(2J_m)}{[1-2n_F(2J_e)][1-2n_F(2J_m)]} = \cos^2\left(\frac{\pi}{2N_\partial}\right). \quad (2.3.46)$$

The geometric mean of the defect densities of the two species, each weighted by $\coth(\beta J_X)$, must equal a constant. In the limit of large boundaries, and when the two coupling constants are well separated, we find that

$$T_c = \frac{4J}{W(J)} \sim \frac{4J}{\ln J} \quad \text{for } J \gg 1, \quad (2.3.47)$$

where $J \equiv J_m/J_e$. We see that, in the limit of large boundaries, the sudden death temperature remains finite. Moreover, T_c for the extended strip partition is approximately twice that of the SPP, as given by (2.3.18). We deduce that multi-particle entanglement is more robust to the effects of thermal fluctuations, but is still killed by an $O(1)$ density of defects.

The corresponding results for the contractible (“open”) strip partition are found using the effective partition function

$$Z = \sum_{\{\tau_s\}} \exp \left\{ -\gamma_{N_\partial} \tau_{N_\partial} - \sum_{s=1}^{N_\partial-1} \gamma_s \tau_s \tau_{s+1} - \sum_{s=1}^{N_\partial} \left[\frac{1}{4} h_s (\tau_s + \tau_{s+1} + 2) + \ln \cosh \beta J_m \right] \right\}. \quad (2.3.48)$$

This Hamiltonian embodies the fact that the plaquette $p = s = N_\partial$ at the end of the boundary is adjacent to only one star. As required, (2.3.48) reduces to the SPP result of Sec. 2.3.1 in the special case $N_\partial \equiv 1$.

2.4 Discussion

The main results of this chapter can be summarised by the complete entanglement “phase diagram” for the two-dimensional toric code, which is presented in Fig. 2.5. At sufficiently low temperatures, $T < J_e$ (recall that, without loss of generality, $J_e \leq J_m$), the density of defects in the system is exponentially suppressed, and the negativity remains undegraded from its zero temperature value. As temperature is increased, we eventually reach $T_{\text{spp}} \sim J_m$ corresponding to the excitation of the higher-energy magnetic vortices, which lead to the sudden death of entanglement between adjacent

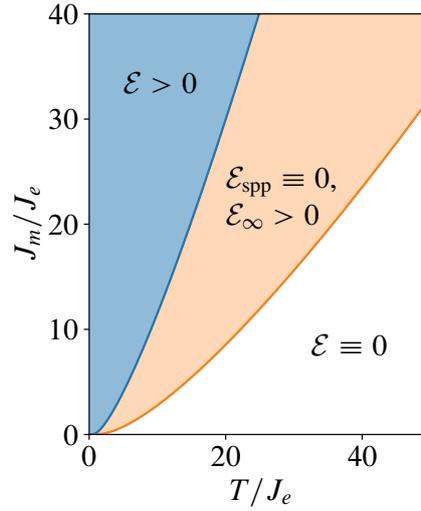


Figure 2.5: The complete entanglement “phase diagram” for the toric code. In the blue region, i.e., $T < T_{\text{spp}}$, all subsystems are entangled with one another (as long as the boundary contains at least one star). In the orange region, $T_{\text{spp}} < T < T_{\infty}$, whether or not two subsystems are entangled depends on their size; the smallest subsystems become separable, while macroscopic regions of the lattice remain entangled. Finally, in the white region, $T > T_{\infty}$, all subsystems exhibit a vanishing negativity, indicative of separability, irrespective of their size.

stars and plaquettes in the lattice. Just above this temperature, SPP subsystems become separable, while subsystems composed of *multiple* stars and plaquettes remain entangled. When temperature is increased further still, larger subsystems also begin to become separable until, finally, for all temperatures $T > T_{\infty} \simeq 2T_{\text{spp}}$, *all* subsystems, irrespective of their size, are separable.

The interpretation of the sudden death temperature for subsystems larger than the SPP is less clear-cut, as a vanishing negativity is only a necessary condition for separability. It represents a lower bound for the temperature above which *all* entanglement vanishes, including any quantum correlations not picked up by the negativity. Any entanglement that may be present in the system above T_c is, however, not able to be distilled, and as a result is often referred to as *bound entanglement* [245]. We would welcome the introduction of an entanglement monotone that is both necessary and sufficient for separability, but expect on physical grounds that our conclusions would remain unchanged up to $O(1)$ factors.

When one of the two coupling constants in the model is infinite, $J_m \rightarrow \infty$, the formation of magnetic vortices is rendered energetically impossible for any finite temperature. Conversely, electric charges are present in the system. This hard constraint means that the subspace of thermally accessible states is substantially smaller than the full Hilbert space, and the calculation is significantly simplified. The hard-constrained model has a substantially simpler expression for the negativity

$$e^{\mathcal{E}(T)} = \frac{\cosh\left[\frac{1}{2}(N - N_{\partial})K_e\right]}{\cosh\left(\frac{1}{2}NK_e\right)} \left[2 \cosh\left(\frac{K_e}{2}\right)\right]^{N_{\partial}}. \quad (2.4.1)$$

In this limit, the negativity does not exhibit any discontinuities. Taking the thermodynamic limit $N \rightarrow \infty$ while keeping the length of the boundary N_∂ finite, the expression (2.4.1) simplifies further, and its extensivity in the number of boundary stars becomes apparent. Introducing the number of non-defective stars $n_h = 1 - n_F(2J_e)$, we find that the negativity becomes

$$\mathcal{E}(T) = N_\partial \ln(2n_h). \quad (2.4.2)$$

In the thermodynamic limit, the hard-constrained model satisfies an exact area law for all temperatures. Numerical results in the regime $J_m < \infty$ for (2.3.35) show that the area law behaviour is present more generally.

The $1/T$ decay of the negativity above $T \sim J_e$ that was observed for the SPP subsystem is no longer cut off at high temperatures. Sending $J_m \rightarrow \infty$ prevents the thermal excitation of magnetic defects, thereby locking in the zero temperature magnetic loop structure. As a result, thermal fluctuations are never sufficient to completely rid the system of all quantum entanglement, which suffers only a slow – power law – demise. This provides a potentially physical example of how a hard constraint applied to a system with a macroscopically degenerate projected manifold can result in thermally robust entanglement. This is likely to be the case for instance in other spin liquids and frustrated magnetic systems in general, when projected down to their low-energy states.

3 | Dynamics of spinons at finite temperature

3.1 Motivation

In the previous chapter, we saw how an $O(1)$ density of thermally excited defects was required to start disrupting the entanglement between two subsystems in direct contact in the toric code. Here, we are interested in the impact of the thermally excited defects on the system's *dynamics*, where the nontrivial statistical angle between the two species of quasiparticle plays an important role.

Experimentally, quantum spin liquid (QSL) candidate materials (for a review, see Refs. [51, 246]) exhibit broad continua in inelastic neutron scattering, indicative of fractionalisation of the emergent quasiparticles (spinons) [247–249]. Although suggestive, such diffuse features are not specific to QSLs, and it is desirable to have more concrete experimental signatures of QSL behaviour. In the following, we will focus on gapped QSLs possessing topological order, characterised by, e.g., subleading corrections to the ground state entanglement entropy [61, 62], and a ground state degeneracy that depends on the genus of the space on which the system resides [250]. While concrete and unambiguous experimental evidence for these unusual ground state properties remains in general unavailable, the exchange statistics of the quasiparticles and their fractional quantum numbers arguably offer some of the most promising routes to unique and experimentally accessible signatures of topological order. We take the stance that, rather than being a hindrance, finite temperature behaviour can in fact offer a number of signatures of fractionalisation and anyonic statistics, and thence of quantum spin liquid behaviour [251–255].

Specifically, we are interested in understanding the role of fractional statistics on the interplay between quasiparticle excitations, in the intermediate temperature range where one species of quasiparticle (visons) is thermally excited and acts as a stochastic background for another species (spinons), which are conversely sparse and hop coherently across the lattice. This is indeed a situation relevant to several realistic Hamiltonians for quantum spin liquids, where there is a large separation between the energy costs of different species of quasiparticle (one can think for example

of quantum spin ice [98], Kitaev materials [50] and valence bond systems [256]). This temperature regime is analogous to the hard constrained limit $J_m \rightarrow \infty$ from the previous chapter, in which magnetic defects are forbidden and the zero temperature magnetic loop structure remains unbroken, while electric charges are excited with a nonzero density.

In this chapter, we derive an effective Hamiltonian for the energetically costly spinons, wherein the visons enter as pointlike sources of magnetic flux, a direct consequence of their mutually semionic statistics. Within a typical vison sector, the effective Hamiltonian therefore maps onto a random flux model, where temperature controls the effective strength of the disorder. The central result of this chapter is an analytical solution within the self-retracing path approximation, in which the effect of the visons is to constrain the worldlines of the spinons to live on a Bethe lattice. Specifically, we derive analytical expressions for the spinon density profile as a function of space and time, which capture the short-time ballistic propagation, the crossover to quantum diffusive behaviour and the non-Gaussian nature of the density profile.

3.2 Hopping Hamiltonian and perturbation theory

In the following, we will consider the behaviour of the \mathbb{Z}_2 lattice gauge theory Hamiltonian introduced in Sec. 1.4.2

$$\hat{H}_{\mathbb{Z}_2} = -J \sum_s \hat{A}_s - h \sum_i \hat{\sigma}_i^z, \quad (3.2.1)$$

where s label the sites of the square lattice, while i indexes the bonds, and $\hat{A}_s = \prod_{i \in s} \hat{\sigma}_i^x$. The model may however also be defined on other two-dimensional lattices, and we will later discuss the case where the spins live instead on a kagome lattice (see Appendix A). However, in contrast to Sec. 1.4.2, the Hilbert space upon which (3.2.1) acts is now the full tensor product space $\mathcal{H} = [\mathbb{C}^2]^{\otimes N}$ over all spins, rather than restricting the physical Hilbert space with a Gauss law constraint. That is, the states $|\{\sigma_i^x\}\rangle$ and $\prod_{j \in p} \hat{\sigma}_j^z |\{\sigma_i^x\}\rangle$ are regarded as physically distinct states (as opposed to being equivalent once projected into the gauge-invariant subspace). Correspondingly, we consider all charge sectors of the theory. When the magnetic field is weak, $h \ll J$, we may treat it as a perturbation. In the limiting case where the magnetic field is absent, $h = 0$, the Hamiltonian is classical in the sense that all terms mutually commute, $[\hat{A}_s, \hat{A}_{s'}] = 0$. This classical Hamiltonian can then be diagonalised by tensor product states in the $\hat{\sigma}_i^x$ basis, $|\{\sigma_i^x\}\rangle$. Each level is however macroscopically degenerate, and so one can equally choose superpositions of such product states, including the eigenstates of the toric code Hamiltonian with well-defined electric and magnetic charges. Up to a constant, the classical Hamiltonian may be written

$$\hat{H}_{\mathbb{Z}_2}^c = 2J \hat{Q}, \quad (3.2.2)$$

where $\hat{Q} = \sum_s \hat{Q}_s = \frac{1}{2} \sum_s (1 - \hat{A}_s)$ counts the total number of defective stars. Each defective star therefore costs an energy $2J$.

Aside: A note on naming conventions and the Gauss law constraint. Suppose that we have a lattice gauge theory in which the Gauss law constraint $\hat{Q}_s |\Psi\rangle = 0$ selects the physical Hilbert space. The operators \hat{Q}_s are mutually commuting, $[\hat{Q}_s, \hat{Q}_r] = 0$, and are conserved $[\hat{H}, \hat{Q}_s] = 0$ under dynamics generated by \hat{H} . Violations of the Gauss law constraint, i.e., $\hat{Q}_s |\Psi\rangle = \rho |\Psi\rangle \neq 0$ are then referred to as electric charges, by analogy with quantum electrodynamics, where electric charges correspond to violations of the familiar solenoidal constraint on the electric field $\nabla \cdot \mathbf{E} = 0$.

In our effective Hamiltonian, we are interested in the limit where magnetic field h is perturbatively small, $h \ll J$, since the model is in its deconfined phase. In the strict limit $J \rightarrow \infty$, the term $-J \sum_s \hat{Q}_s$ can be regarded as a hard constraint which enforces that the physically permissible (finite energy) states $|\Psi\rangle$ in the Hilbert space satisfy $\hat{Q}_s |\Psi\rangle = 0$. If J is instead large but finite, then states satisfying $\hat{Q}_s |\Psi\rangle = \rho |\Psi\rangle \neq 0$ are permitted, but energetically costly. We refer to violations of this constraint as electric charges or spinons. Note that this terminology is the opposite of that used in Sec. 1.4.2, where violations of the constraint $\hat{B}_p |\Psi\rangle = |\Psi\rangle$ were referred to as electric charges.

Our choice of terminology comes from the parallels between our Hamiltonian and quantum spin ice. There, as shown in Sec. 1.4.4, the large classical spin ice term in the Hamiltonian selects the ground state sector, which is spanned by the states that contain no monopoles, i.e., $\hat{Q}_t |\Psi\rangle = 0$ for each tetrahedron t . The effective Hamiltonian within the degenerate ground state manifold can then be described using an emergent compact U(1) lattice gauge theory. Violations of the Gauss law constraint (the “electric charges”) are now the monopoles of classical spin ice, which are also commonly referred to as spinons. Therefore, in both cases, we will consistently refer to violations of the large, macroscopically degenerate, classical Hamiltonian as spinons. Similarly, we will refer to excitations within the ground state sector (i.e., the spinon vacuum) as visons. The gapless photon has no analogue in gapped \mathbb{Z}_2 lattice gauge theories.

Recall that the operators $\hat{B}_p = \prod_{i \in p} \hat{\sigma}_i^z$, where p labels the plaquettes, are *exact* conserved quantities for the Hamiltonian (3.2.1), $[\hat{B}_p, \hat{H}_{\mathbb{Z}_2}] = 0$. It is therefore convenient to work with a basis of states that diagonalise the \hat{B}_p operators. Defective plaquettes satisfy $\hat{B}_p |\Psi\rangle = -|\Psi\rangle$, and will be referred to as visons (see the note above regarding our choice of naming conventions). In this basis, the Hamiltonian becomes block diagonal, which allows us to treat each sector containing a distinct configuration of visons separately. Their status as conserved quantities ensures that vison configurations remain precisely static under time evolution generated by (3.2.1), i.e., they behave as quasiparticles with infinite mass. Unlike the toric code Hamiltonian, the \hat{A}_s operators are however not good quantum numbers, $[\hat{A}_s, \hat{H}_{\mathbb{Z}_2}] \neq 0$. Instead, a nonzero magnetic field endows the spinons with dynamics.

To make this picture more quantitative, the magnetic field term in (3.2.1) can alternatively be

written in terms of operators \hat{T}_n that modify a given configuration of spinons

$$\hat{H}_{\mathbb{Z}_2} = 2J\hat{Q} - h\left[\hat{T}_0 + \hat{T}_2 + \hat{T}_{-2}\right]. \quad (3.2.3)$$

Specifically, the operators \hat{T}_n increment the number of spinons by n , i.e., $[\hat{Q}, \hat{T}_n] = n\hat{T}_n$. If we are working with periodic boundary conditions, then spinons can only be created in pairs since $\prod_s \hat{A}_s = \mathbb{1}$. In the presence of open boundary conditions, this condition is relaxed, since the boundary can become a source or sink for single spinons, but this does not modify the bulk Hamiltonian (3.2.3), where spinons are *locally* created or removed in pairs. These operators can be written out explicitly in terms of the spins as

$$\hat{T}_0 = \sum_{\langle ss' \rangle} \left[\frac{1}{2}(\mathbb{1} - \hat{A}_s)\hat{\sigma}_{ss'}^z \frac{1}{2}(\mathbb{1} - \hat{A}_{s'}) + \frac{1}{2}(\mathbb{1} + \hat{A}_s)\hat{\sigma}_{ss'}^z \frac{1}{2}(\mathbb{1} + \hat{A}_{s'}) \right] \quad (3.2.4)$$

$$\hat{T}_2 = \sum_{\langle ss' \rangle} \frac{1}{2}(\mathbb{1} - \hat{A}_s)\hat{\sigma}_{ss'}^z \frac{1}{2}(\mathbb{1} + \hat{A}_{s'}) \quad (3.2.5)$$

$$\hat{T}_{-2} = \sum_{\langle ss' \rangle} \frac{1}{2}(\mathbb{1} + \hat{A}_s)\hat{\sigma}_{ss'}^z \frac{1}{2}(\mathbb{1} - \hat{A}_{s'}), \quad (3.2.6)$$

where $\hat{\sigma}_{ss'}^z$ corresponds to the spin on the bond between adjacent sites s and s' ($\hat{\sigma}_{ss'}^z$ is not a directed variable in the sense that $\hat{\sigma}_{ss'}^z = \hat{\sigma}_{s's}^z$), and we have used $\{\hat{A}_s, \hat{\sigma}_{ss'}^z\} = \{\hat{A}_{s'}, \hat{\sigma}_{ss'}^z\} = 0$. The action of these operators on eigenstates of the toric code is shown schematically in Fig. 3.1. \hat{T}_0 conserves the number of spinons, hopping the quasiparticle to each of the four adjacent sites that are unoccupied. Meanwhile, $\hat{T}_{\pm 2}$ creates (destroys) two spinons residing on adjacent sites, if both sites are unoccupied (occupied). We are now in a position to construct effective Hamiltonians that conserve the number of spinons up to a particular order in the magnetic field.

3.2.1 Perturbation theory and effective Hamiltonians

In order to construct spinon-conserving effective Hamiltonians, we make use of a Schrieffer–Wolff transformation [257, 258], parametrised in terms of a Hermitian operator $\hat{S} = \hat{S}^\dagger$

$$\hat{H}' = e^{i\hat{S}}\hat{H}e^{-i\hat{S}} = \hat{H} + [i\hat{S}, \hat{H}] + \frac{1}{2!}[i\hat{S}, [i\hat{S}, \hat{H}]] + \dots \quad (3.2.7)$$

While the bare Hamiltonian \hat{H} does not conserve the number of spinons, we are able to construct the operator \hat{S} to be of the form $\hat{S} = \sum_k \hat{S}^{(k)}$, where $S^{(k)}$ is chosen such that \hat{H}' conserves the number of spinon quasiparticles up to order $(h/J)^k$. The benefit of proceeding in this manner is that the number of basis states that are connected by H' does not proliferate with successive applications of the Hamiltonian. Equivalently, since the number of quasiparticles is conserved,

we can work in a restricted subspace spanned by states with equal quasiparticle number. This process of eliminating terms that violate spinon number conservation order by order is equivalent to performing standard (degenerate) Rayleigh–Schrödinger perturbation theory in the magnetic field. Since the two Hamiltonians \hat{H}' and \hat{H} are related via a unitary transformation, they are spectrally equivalent. Meanwhile, the eigenstates $|\psi_k\rangle$ of \hat{H} are related to those of \hat{H}' by $|\psi_k\rangle = e^{-i\hat{S}} |\psi'_k\rangle$. The operator $e^{-i\hat{S}}$ can be thought of as “dressing” the eigenstates $|\psi'_k\rangle$ of the number-conserving Hamiltonian. For example, at first order in the magnetic field,

$$i\hat{S} = i\hat{S}^{(1)} = -\frac{\hbar}{4J}(\hat{T}_2 - \hat{T}_{-2}). \quad (3.2.8)$$

Since $\hat{T}_n^\dagger = \hat{T}_{-n}$, the operator $i\hat{S}$ is anti-Hermitian, as it must be. We therefore observe that $e^{-i\hat{S}}$ dresses the eigenstates of \hat{H}' with quasiparticle pairs. Substituting this expression into the BCH expansion (3.2.7), we deduce that $[iS^{(1)}, H_{\mathbb{Z}_2}] = \hbar(T_2 + T_{-2})$, which by construction cancels the first order terms in \hat{H} that do not conserve particle number. This allows us to deduce the *second* order contribution to the transformed Hamiltonian

$$\hat{H}' = 2J\hat{Q} - \hbar\hat{T}_0 + \frac{\hbar^2}{4J}[\hat{T}_2, \hat{T}_{-2}] + \dots \quad (3.2.9)$$

The term $\propto \hbar$ describes hopping of the spinons to adjacent lattice sites. The second term $\propto \hbar^2$ encapsulates two processes: virtual creation and annihilation of spinon pairs, and next-nearest neighbour hopping of spinons. This process can, in principle, be carried out to arbitrarily high orders in the magnetic field \hbar . Methods such as perturbative continuous unitary transformations (pCUT) can be used to efficiently carry out such high order expansions [259–261]. Indeed, this method has been used to calculate the quasiparticle gap and dispersion in the toric code perturbed by a uniform field, from which one can infer the boundaries of the deconfined phase [65, 66]. In this thesis we will not go beyond fourth order in the magnetic field, as the effective Hamiltonian at this order already contains the relevant physics. For concreteness, at third order in the applied field,

$$\hat{H}' = 2J\hat{Q} - \hbar\hat{T}_0 + \frac{\hbar^2}{4J}[\hat{T}_2, \hat{T}_{-2}] - \frac{\hbar^3}{32J^2} \left\{ [\hat{T}_2, [\hat{T}_0, \hat{T}_{-2}]] + [\hat{T}_{-2}, [\hat{T}_0, \hat{T}_2]] \right\} + \dots \quad (3.2.10)$$

3.2.2 Spinon vacuum

The spinon vacuum (spanned by the classical ground states that contain strictly zero spinons) is a special case since there are no spinons that the operator \hat{T}_0 can hop, i.e., all states $|\Psi\rangle$ belonging to the spinon vacuum satisfy $\hat{T}_0 |\Psi\rangle = 0$. The first contribution therefore comes at second order, where the term $\hat{T}_{-2}\hat{T}_2$ leads to the creation and annihilation of neighbouring spinon pairs. Such a

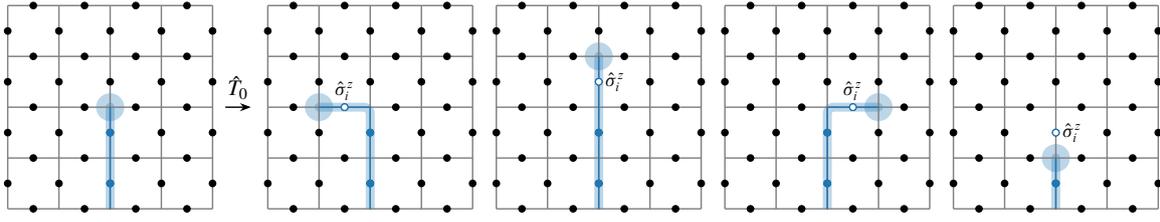


Figure 3.1: The action of the number-conserving operator \hat{T}_0 on an isolated spinon. The operator \hat{T}_0 hops the spinon to each of the adjacent four sites. This process is only possible if the adjacent sites are unoccupied.

term however does not distinguish between different vison configurations, shifting all states with the same vison configuration identically. In the effective Hamiltonian, this leads to a trivial energy offset. The first order at which the macroscopic degeneracy of the spinon vacuum is lifted is at *fourth* order in the magnetic field due to “ring exchange” processes. Extending Eq. (3.2.9) to fourth order, the effective Hamiltonian within the spinon vacuum is

$$\hat{H}_{\text{eff}}^{(0)} = -J \sum_s \hat{A}_s - \frac{5}{16} \frac{h^4}{J^3} \sum_p \hat{B}_p + \text{const.}, \quad (3.2.11)$$

where the constant energy offset arises from the virtual creation and annihilation of spinon pairs. We recognise (3.2.11) as the toric code Hamiltonian from Chapter 2. Higher order corrections to the above Hamiltonian do not change the eigenstates of the effective Hamiltonian, $\{|B_p\rangle\}$, but will modify their energies, with short-ranged interactions between the \hat{B}_p operators being generated at 8th order and above in the magnetic field.

3.2.3 Effective Hamiltonian for an isolated spinon

In order to derive an effective Hamiltonian for the spinons, we make a “gauge choice” by fixing the string $\hat{S}(\gamma_i)$ associated with a spinon residing on site i . We imagine taking the other end of the string to infinity, which may end on a defect or the boundary of the system, for example. In the absence of boundaries or defects, the other end of the string must correspond to another quasiparticle. The string $\hat{S}(\gamma_i) = \prod_{(ss') \in \gamma_i} \hat{\sigma}_{ss'}^z$ depends on the path γ_i , which terminates on site i . Having made this choice, we can calculate the matrix elements of the effective Hamiltonian (3.2.9). For neighbouring sites s and s'

$$t_{ss'} \equiv -h \langle \{B_p\}, s' | \hat{\sigma}_{ss'}^z | \{B_p\}, s \rangle = -h \langle \{B_p\} | \hat{S}(\gamma_{s'}) \hat{\sigma}_{ss'}^z \hat{S}(\gamma_s) | \{B_p\} \rangle. \quad (3.2.12)$$

Now, the product $\hat{S}(\gamma_{s'}) \hat{\sigma}_{ss'}^z \hat{S}(\gamma_s)$ forms a closed loop $C_{ss'}$, whose boundary is given by $\gamma_s \cup \gamma_{s'}$ and the bond connecting s and s' . The product of spins on the loop can be extended to the product of \hat{B}_p operators on the interior of the region bounded by $C_{ss'}$: $\prod_{i \in C_{ss'}} \hat{\sigma}_i^z = \prod_{p \in R_{ss'}} \hat{B}_p$, where

$\partial R_{ss'} = C_{ss'}$. Therefore Eq. (3.2.12) reduces to

$$t_{ss'} = -h \prod_{p \in R_{ss'}} B_p. \quad (3.2.13)$$

Clearly the hopping amplitudes depend explicitly on the choice of the string operators (equivalently, the paths γ_i). Note that choosing a *different* string $\hat{S}(\gamma'_i)$ does not change the state, only its phase, since $|\langle \{B_p\} | \hat{S}(\gamma_i) \hat{S}(\gamma'_i) | \{B_p\} \rangle|^2 = 1$. Moving the spinon around a plaquette p in a closed loop, we find that

$$\prod_{\langle ss' \rangle \in p} t_{ss'} = h^4 B_p, \quad (3.2.14)$$

independent of the choice of string. This result follows from the fact that $\hat{S}(\gamma_i)^2 = \mathbb{1}$. Therefore the visons act as sources of emergent magnetic flux for the spinons; a consequence of their mutually semionic statistics. The eigenvalues B_p may be written in the form $B_p = e^{i\phi_p}$, with $\phi_p \in \{0, \pi\}$, i.e., the visons act as point-like sources of flux, with a magnitude equal to half the flux quantum. Since only gauge-invariant quantities are accessible in experiment, the choice of string γ_i is essentially immaterial.

Therefore, at first order in the magnetic field, the effective Hamiltonian in a typical vison sector corresponds to a tight-binding model with “disorder” in the hopping amplitudes

$$\hat{H}_{\text{eff}} = -h \sum_{\langle ss' \rangle} \left(\hat{b}_s^\dagger e^{i\phi_{ss'}} \hat{b}_{s'} + \text{H.c.} \right), \quad (3.2.15)$$

where the Peierls phases $\phi_{ss'}$ are determined by the location of the visons, and the operators $\hat{b}_s, \hat{b}_s^\dagger$ represent hardcore bosons. In principle, this process can be extended to arbitrary order in the magnetic field. A particularly convenient gauge choice when using open boundary conditions is one in which all vertical bonds have $\phi_{ss'} = 0$, while in the horizontal direction the hopping amplitudes change sign whenever a π flux is encountered.

The motion of spinons in our model is therefore equivalent to a quantum particle propagating through a background of randomly-placed, static* \mathbb{Z}_2 fluxes. We also briefly consider for comparison the generic cases of other values of the threaded fluxes, $2\pi/n$, $n = 3, 4, \dots$, and, in particular, the limiting case of the so-called continuous flux model in which the fluxes ϕ threading the plaquettes are drawn from the uniform distribution over $\phi \in [0, 2\pi)$. Analogously to the random \mathbb{Z}_2 flux model, the model with random $2\pi/n$ fluxes arises from a \mathbb{Z}_n lattice gauge theory in a similar finite temperature regime where the flux excitations are thermally populated whilst the elementary charge

*Introducing a nonzero projection of the magnetic field onto the x axis, h_x , gives rise to vison dynamics on a timescale h_x^{-1} . However, provided that $h_x \ll h$, over the intermediate timescales of interest to the motion of the spinons, $h^{-1} \ll h_x^{-1}$, the visons may be treated as static and our results apply.

excitations remain coherent. Likewise, the continuous flux model describes the motion of a charged particle through an incoherent U(1) gauge field.

At this point, we take a step back and look at what we have accomplished by performing the Schrieffer–Wolff transformation. The number of bare spinons is not conserved by $\hat{H}_{\mathbb{Z}_2}$, but it is possible to transform into a basis in which spinon quasiparticles are conserved up to a particular order in the small parameter h/J . In the spinon vacuum sector, the effective Hamiltonian has the same eigenstates as the toric code (to all orders in perturbation theory), while in the first excited sector, containing a single spinon, the model reduces to a tight binding model of bosonic spinons hopping in an emergent magnetic field. The visons act as sources of π flux for the spinons, just as the plaquette operators \hat{W}_p acted as sources of π flux for the Majorana fermions in the honeycomb model in Sec. 1.4.5. One should however keep in mind that (3.2.15) is obtained within perturbation theory, and its eigenstates are related to those of the original Hamiltonian $\hat{H}_{\mathbb{Z}_2}$ by the unitary dressing operator $e^{-i\hat{S}}$ [given to lowest order in Eq. (3.2.8)], which decorates the states with quasiparticle pairs.

Since the lattice is bipartite, the first order effective Hamiltonian exhibits a chiral symmetry. That is, if we refer to the two sublattices as A and B , then A is connected *only* to B and vice versa. Further, the hopping amplitudes can be chosen to be entirely real; a consequence of time reversal symmetry. This implies that the Hamiltonian may be written (after a suitable ordering of basis states) as

$$H = \begin{pmatrix} 0 & H_{AB} \\ H_{AB}^\dagger & 0 \end{pmatrix}. \quad (3.2.16)$$

These features place the Hamiltonian (3.2.15) in a different universality class to “conventional” Anderson localisation in two dimensions, falling instead under the heading of bipartite random hopping models [262–264]. As discussed in Sec. 1.5.1, the pure off-diagonal nature of the disorder leads to a strongly divergent density of states near zero energy. It must be remembered however that the chiral symmetry is not an exact symmetry, neither of the bare Hamiltonian, nor of full effective Hamiltonian. Already at second order in the magnetic field the A sublattice is connected to itself, and similarly for B . Although the chiral symmetry is not exact, it is approximate for sufficiently small magnetic fields and will dominate the behaviour of the system over some time scale (the effect of chiral symmetry breaking terms modify substantially the low-energy properties of the system over some small energy window around $E = 0$, which manifest at correspondingly long times).

3.2.4 Finite temperature

At zero temperature, the ground state of the system is determined by the lowest energy state in the spinon vacuum sector, i.e., the ground state of (3.2.11). As temperature is raised from zero, there

are a number of different regimes to consider. First, at temperatures

$$T \ll \frac{h^4}{J^3} \ll h \ll J, \quad (3.2.17)$$

both spinons and visons are exponentially suppressed in temperature by their gaps, and the equilibrium behaviour of the system is not qualitatively distinct from zero temperature. If a spinon is artificially injected into the system, then the effective Hamiltonian corresponds to a clean (flux free) tight binding model*. As temperature is raised further,

$$\frac{h^4}{J^3} \sim T \ll h \ll J, \quad (3.2.18)$$

visons are excited with a finite density, and the coherence of the spinon vacuum is partially degraded. The finite vison density manifests in the effective Hamiltonian of an isolated spinon as off-diagonal disorder, i.e., disorder in hopping amplitudes. As temperature is raised further still

$$\frac{h^4}{J^3} \ll T \ll h \ll J, \quad (3.2.19)$$

the density of visons $n_v \simeq 1/2$, and the quantum coherence of the spinon vacuum is now completely disrupted. The density matrix of the system within this temperature regime may be approximated by an incoherent mixture of all states belonging to the spinon vacuum

$$\hat{\rho} = \sum_{\{B_p\}} |\{B_p\}\rangle \langle \{B_p\}| = \sum_{\{\sigma_i^x\}: A_s=1} |\{\sigma_i^x\}\rangle \langle \{\sigma_i^x\}|. \quad (3.2.20)$$

The density matrix now behaves as a projector into the spinon vacuum and subjects spinons to disordered flux configurations drawn from an infinite temperature distribution. Since $T \ll J$, the equilibrium density of spinon quasiparticles is still exponentially suppressed†.

*As we will show later, the limits of zero temperature and infinite time do not commute. The behaviour is not qualitatively distinct from zero temperature up to a time scale set by the characteristic separation of visons, $t \lesssim n_v^{-1}$.

†We refer to the eigenstates of the effective Hamiltonian \hat{H}' as spinon quasiparticles. The number of *bare* spinons is set not only by the ratio of T/J , but also by h/J , since even the ground state contains a finite density of bare spinons set by $O(h^2/J^2)$.

3.3 Lattice walks

3.3.1 Single-particle Green's function

One quantity of interest is the single-spinon* on-site Green's function for the effective Hamiltonian (3.2.15), defined by

$$G_{ii}(t) = \langle\langle \hat{b}_i(t) \hat{b}_i^\dagger(0) \rangle\rangle, \quad (3.3.1)$$

where the double angled brackets $\langle\langle \dots \rangle\rangle$ refer to both the quantum expectation value and thermal averaging, which manifests as an average over vison (flux) configurations. The on-site Green's function gives us access to the finite-temperature single-particle density of states $\rho(\omega)$ for spinons. The form of $G_{ii}(\omega)$ within the self-retracing path approximation is well known in the context of the t - J_z model [265–268]. We include its derivation using the method of generating functions for the sake of completeness. Our methodology may then be generalised to determine the off-diagonal elements of the Green's function, $G_{ij}(t)$, with $i \neq j$. These quantities are not on their own gauge invariant and must be multiplied by the phases corresponding to a given lattice path connecting sites i and j in order to construct a gauge invariant quantity [269].

Formally expanding $e^{-i\hat{H}t}$ governing the time evolution in (3.3.1) as a power series in time[†], we are able to convert the problem into a summation over discrete lattice paths γ , where the particle moves one lattice spacing per step. Integrating out the vison configurations, we arrive at

$$G_{ii}(t) = \sum_{\ell=0}^{\infty} \frac{(iht)^\ell}{\ell!} \sum_{\gamma \in \Gamma(\ell)} e^{-A_\gamma/\xi^2(T)} \quad (3.3.2)$$

$$= \sum_{\ell=0}^{\infty} \frac{(iht)^\ell}{\ell!} |\Gamma(\ell)| \left\langle e^{-A_\gamma/\xi^2(T)} \right\rangle_{\gamma \in \Gamma(\ell)}, \quad (3.3.3)$$

where $\Gamma(\ell)$ is the set of all paths of length ℓ that begin and end at the site i , and $A_\gamma = \sum_p A_p(\gamma)$ is the “area” enclosed by γ : each plaquette p contributes an area $A_p(\gamma) = [1 - (-1)^{w_p(\gamma)}]/2$ if it is encircled a total of $w_p(\gamma)$ times by the path γ . In the continuous flux model, a given plaquette contributes only if $w_p(\gamma) = 0$. In the second line, we have written the ℓ th term as an average over all random walks that return to their starting point after ℓ steps.

The length scale $\xi(T)$ appearing in the exponential,

$$\xi^2(T) = \frac{1}{-\ln \tanh[5\beta h^4/(16J^3)]}, \quad (3.3.4)$$

*On account of their fractionalised nature, spinons are created in pairs. The single-particle properties are relevant if treating the two spinons as independent.

[†]One may alternatively calculate $G_{ii}(\omega)$ as a power series in $1/\omega$ using similar methods, as in Ref. [265].

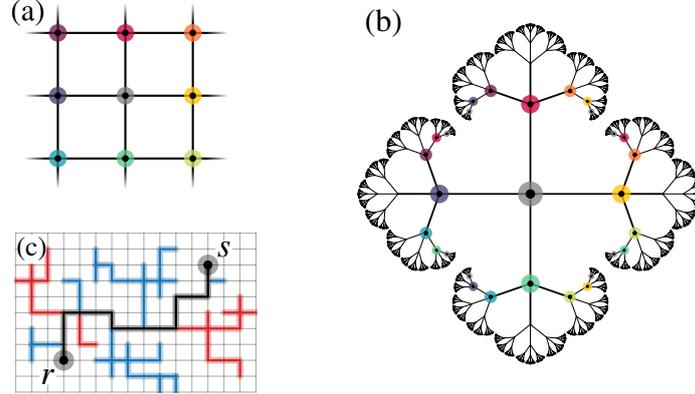


Figure 3.2: Mapping from the square lattice, (a), to the Bethe lattice with coordination number $z = 4$, (b), used for the calculation of the spinon density profile. Each site on the square lattice is mapped onto multiple sites on the Bethe lattice, as indicated by the coloured circles. An example of a perfectly self-retracing round trip ($r \rightarrow s \rightarrow r$) on the square lattice is shown in (c). The nonreversing base path connecting r and s is represented by the thick black line, while the self-retracing excursions that decorate the base path on the outward (return) trip are shown as thick blue (red) lines. Such a walk encloses precisely zero area and contributes to the high-temperature expansion of the transition probability $P_{r \rightarrow s}(t)$.

corresponds approximately to the average distance between visons, $\sim n_v^{-1/2}$, in the dilute-vison limit, where $n_v \sim e^{-10\beta h^4/16J^3}$. Note that $\xi \rightarrow 0^+$ for high temperatures, corresponding to the vison-dense limit, $n_v \rightarrow 1/2$. At *any* nonzero temperature, paths that enclose a large area with respect to ξ^2 are exponentially suppressed*, a manifestation of the Aharonov–Bohm effect. Recall that the zero-temperature ($n_v = 0$) dynamics of the spinon is equivalent to a free quantum particle at all times t . The limits of infinite time and zero temperature therefore do not commute.

Aside: Ordinary versus exponential. Given an ordinary generating function

$$f(x) = \sum_n f_n x^n, \quad (3.3.5)$$

we would like to be able to derive the corresponding *exponential* generating function, defined by the expression

$$\mathcal{F}(x) = \sum_n f_n \frac{x^n}{n!}. \quad (3.3.6)$$

This can be accomplished most neatly using complex contour integration with the formula

$$\frac{x^n}{n!} = \frac{1}{2\pi i} \int_C dz \frac{e^{xz}}{z^{n+1}}, \quad (3.3.7)$$

*This may be viewed as an analogue of the high-temperature area law for loop correlation functions $\langle \prod_{\ell \in C} \sigma^z(\ell) \rangle \sim \exp(-A)$ in lattice gauge theory [270].

where the contour C encircles the origin. In fact, the contour cannot be chosen arbitrarily since $f(x)$ may have a finite radius of convergence. Substituting (3.3.7) into the expression for the exponential generating function, we observe that

$$\mathcal{F}(x) = \frac{1}{2\pi i} \sum_n f_n \int_C dz \frac{e^{xz}}{z^{n+1}} = \frac{1}{2\pi i} \int_C dz \frac{e^{xz}}{z} f\left(\frac{1}{z}\right). \quad (3.3.8)$$

We therefore deduce that C must be chosen such that $1/z$ lies within the radius of convergence. Equation (3.3.8) provides us with a relationship between the ordinary and exponential generating functions.

From the ordinary generating function $g(x; a) = \sum_{n,m} g_{nm} x^n a^m$ for walks $\gamma \in \Gamma(\ell)$, where the generating variables x and a are associated with path length n and area enclosed m , respectively, one can observe that $G_{ii}(t)$ is equal to the corresponding exponential generating function $\tilde{g}(x; a) \equiv \sum_{n,m} g_{nm} x^n a^m / n!$ via

$$G_{ii}(t) = \tilde{g}(iht; e^{-1/\xi^2}). \quad (3.3.9)$$

The effect of changing temperature is to alter the relative weight of the different lattice walks, classified according to the area that they enclose. At zero temperature, all paths of a given length contribute with equal weight, while at infinite temperature only those paths that enclose precisely zero area contribute. We expect that the latter result describes the limiting behaviour for high temperatures, $T \gg h^4/J^3$.

3.3.2 Density evolution

Motivated by the study of finite temperature dynamical spin–spin correlators, we would like to quantify the propagation of a pair of spinons after being created locally on adjacent sites. As a first approximation, we solve the single-particle problem, which we are able to treat analytically. In particular, we calculate (i) the site-resolved density profile for short times, $ht = O(1)$, accessible for instance in ultracold atomic experiments, and (ii) the asymptotic moments of the density distribution, $\langle \mathbf{r}^{2k}(t) \rangle$, which characterise the behaviour of the spinon profile over a large range of intermediate time scales.

The (gauge invariant) probability for the spinon to move from site 0 to site s in a time t in the presence of a given vison (flux) configuration $\{\phi_p\}$ is given by

$$P_s(\{\phi_p\}; t) = \langle \{\phi_p\} | \hat{b}_0^\dagger(0) \hat{b}_s^\dagger(t) \hat{b}_s(t) \hat{b}_0^\dagger(0) | \{\phi_p\} \rangle \quad (3.3.10)$$

$$= \left| \langle \{\phi_p\} | \hat{b}_s e^{-i\hat{H}t} \hat{b}_0^\dagger | \{\phi_p\} \rangle \right|^2. \quad (3.3.11)$$

In a similar manner to $G_{ii}(t)$, we are able to write the transition probability $P_s(t)$ in terms of summation over outward (γ) and return (γ') lattice paths. After integrating over the possible flux

configurations $\{\phi_p\}$ with the appropriate Boltzmann weight, the probability reads

$$P_s(t) = \sum_{\ell, \ell'=0}^{\infty} (-1)^{\ell} \frac{(iht)^{\ell+\ell'}}{\ell! \ell'!} \sum_{\substack{\gamma \in \Gamma_s(\ell) \\ \gamma' \in \Gamma_s(\ell')}} e^{-A_{\gamma \cup \gamma'} / \xi^2}, \quad (3.3.12)$$

where $\Gamma_s(\ell)$ is the set of all paths of length ℓ that connect the sites 0 and s , and $A_{\gamma \cup \gamma'}$ is the area enclosed by the closed path $\gamma \cup \gamma'$. Knowledge of $P_s(t)$ for all sites s gives us complete information about the density distribution $\rho(\mathbf{r}, t)$ as a function of time.

3.3.3 Interpretation

We have shown that in both instances the problem of determining single-spinon motion in a sea of thermally excited visons may be mapped onto the combinatorial problem of enumerating discrete lattice walks. At precisely zero temperature, the system is free of vison excitations, $\xi = \infty$, and all paths of a given length contribute with equal weight*. In this limit, the effective Hamiltonian is simply a two-dimensional tight-binding model with nearest neighbour hopping, and the spinon propagates ballistically. Conversely, at temperatures which are high with respect to the energy cost for vison creation, $T > h^4/J^3$, the hopping amplitudes are maximally disordered, and only walks that enclose exactly “zero area” (as defined previously) contribute. (Note that the notion of zero enclosed area trivially extends to the case of fluxes of magnitude $2\pi m/n$, with $m, n \in \mathbb{N}$, threading the plaquettes.) We focus primarily on this high-temperature limit in order to contrast with the known behaviour at $T = 0$.

Exact enumeration of all such zero-area paths on a generic lattice with coordination number z is a very tall order. To tackle this problem, we discuss a limit where the problem becomes analytically tractable. Namely, we consider perfectly self-retracing paths [265], which necessarily satisfy $A_\gamma = 0$. As we shall see, this is a particularly relevant subset of walks the smaller the threaded fluxes are, and particularly for the continuous U(1) flux case. For comparison, in Sec. 3.6, we also solve numerically for the time evolution generated by effective Hamiltonians of the form (3.2.15) using a high order Suzuki–Trotter decomposition [271]. We will also discuss the lower temperature regime $T \lesssim h^4/J^3$ in the context of these simulations.

3.4 Self-retracing paths: Generating functions

A perfectly self-retracing path corresponds to a lattice walk with no closed cycles in which every link on the ‘outwards’ path is retraced in the opposite direction on the ‘return’ path. More precisely,

*Since the limits of infinite time and zero temperature do not commute. We are here referring to taking $T \rightarrow 0$ before $t \rightarrow \infty$.

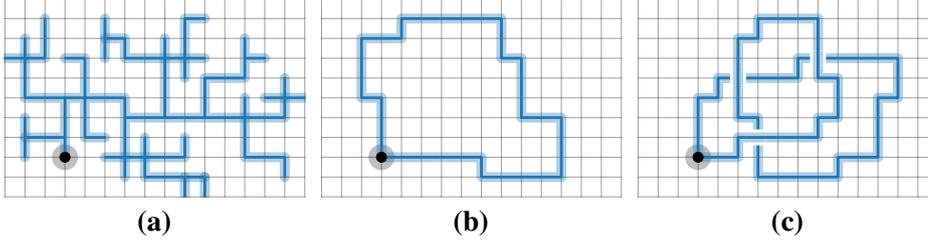


Figure 3.3: Three examples of closed lattice walks, γ , which begin and end on the black circle. A perfectly self-retracing path – the only type of path included in the Bethe lattice mapping – is shown in (a). In (b) and (c) the walk includes closed cycles, which have the potential to be non-self-retracing. In general, a walk γ contributes to the lattice path expansion at high temperatures if $\langle \exp(i \sum_{(ij) \in \gamma} \phi_{ij}) \rangle_{\{\phi_{ij}\}} = 1$. In the case of π -fluxes threading the plaquettes, a walk in which the loop (b) is traversed an even number of times in the *same* direction leads to a nonzero contribution, $\langle e^{2ni\phi} \rangle_{\mathbb{Z}_2} = 1$. Such a path with winding number $w = 2n$ ($n \in \mathbb{Z}_{\neq 0}$) is *not* self-retracing, and so is not captured by the Bethe lattice mapping. For continuous fluxes, however, such non-self-retracing paths of the form (b) with nonzero winding number $w = 2n$ are not present in the lattice path expansion after averaging over flux configurations, since $\langle e^{2ni\phi} \rangle_{U(1)} = 0$. Paths that self-intersect multiple times, as in (c), can be traversed in multiple ways in the reverse direction, only one of which is self-retracing. All other paths are not accounted for by the Bethe lattice mapping, whereas they do however contribute to the continuous flux case. For these reasons, we expect the Bethe lattice mapping to better approximate the continuous flux model where a significantly larger fraction of permitted lattice walks are correctly enumerated.

there is a one-to-one mapping between self-retracing paths on a lattice, \mathcal{L}_z , with coordination number z , and closed walks on a Bethe lattice, \mathcal{B}_z , with branching ratio $z - 1$, as shown for the case $z = 4$ in Fig. 3.2.

Fractal lattices in general are a useful tool for obtaining exact solutions, and have recently been used in a similar context to calculate the spectrum of itinerant excitations in quantum spin ice at zero temperature [272, 273], where gauge field effects lead to a configuration space which is well-approximated by the Husimi cactus graph. Examples of perfectly self-retracing walks, and classes of walks which are *not* captured by the self-retracing path approximation are shown in Fig. 3.3. The latter are analogous in spirit to the Trugman path [274] in the context of single hole propagation in the t - J model. Since the paths that are not accounted for only become relevant at long times, one may expect that this approximation works well for the dynamics of the system at the intermediate timescales of interest, at least for the continuous flux model.

In this section we will derive the generating functions for walks on a Bethe lattice with branching ratio $z - 1$, which are necessary to describe analytically the form of the density profile in the limit of high temperature (high flux density).

3.4.1 Closed walks

We first consider the ordinary generating function $T_0^{(z)}(x) = \sum_n t_n^{(z)} x^n$ for closed walks on a Bethe lattice \mathcal{B}_z with branching ratio $z - 1$, where, by definition, $t_n^{(z)}$ is the number of closed walks that begin and end at the same site, which may be used to define the root node (or origin) of the Bethe lattice. This generating function is directly related to the single-particle density of states. Note that the lack of closed cycles implies that all closed walks on \mathcal{B}_z are necessarily self-retracing, and further permits the following decomposition of the generating function

$$T_0^{(z)}(x) = 1 + zx^2 T_0^{(z)}(x) T_1^{(z)}(x), \quad (3.4.1)$$

where $T_k^{(z)}(x)$ is the generating function for walks beginning and ending at a depth of k on the lattice (always remaining at a depth $\geq k$). This is because any (self-retracing) path can be decomposed as

- (i) the trivial walk,
- (ii) (a) hopping to one of the z nearest neighbours,
 - (b) performing a self-retracing walk that begins and ends at depth $k = 1$,
 - (c) hopping back to the origin,
 - (d) performing a self-retracing walk that begins and ends at the origin.

A similar argument can be made for all subsequent depths with $k \geq 1$, such that the generating functions decompose as

$$T_k^{(z)}(x) = 1 + (z - 1)x^2 T_k^{(z)}(x) T_{k+1}^{(z)}(x). \quad (3.4.2)$$

One can therefore express the original generating function $T_0^{(z)}$, for paths beginning and ending at the origin, as an infinite continued fraction

$$T_0^{(z)}(x) = \frac{1}{1 - \frac{zx^2}{1 - \frac{(z-1)x^2}{1 - \frac{(z-1)x^2}{1 - \dots}}}}. \quad (3.4.3)$$

The self-similar nature of \mathcal{B}_z implies that $T_k^{(z)}(x) = T_{k+1}^{(z)}(x)$ for $k \geq 1$ (on an infinite lattice), and the continued fraction can be written in closed form (choosing the sign in front of the square root

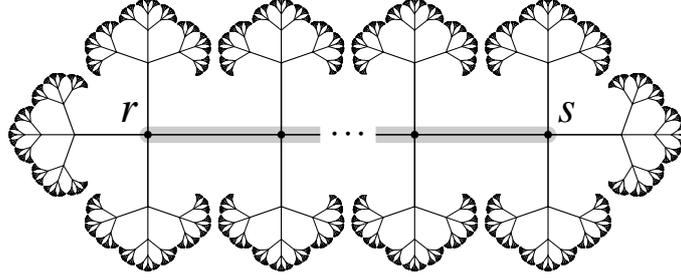


Figure 3.4: Example of a path from r to s on the Bethe lattice \mathcal{B}_4 , which maps to a *nonreversing* walk of length ℓ on the square lattice. For any given two sites on the square lattice, there are multiple nonreversing walks that connect the two sites, enumerated by the generating function $C_s(x)$.

such that $T_0^{(z)}(x) \rightarrow 1$ as $x \rightarrow 0$)

$$T_0^{(z)}(x) = \frac{2(z-1)}{z-2+z\sqrt{1-4(z-1)x^2}}, \quad (3.4.4)$$

consistent with, e.g., the results of Ref. [275]. For $z = 4$ (corresponding to the square lattice at high temperatures), this expression evaluates to

$$T_0^{(4)}(x) = \frac{3}{1+2\sqrt{1-12x^2}} = 1 + 4x^2 + 28x^4 + \dots \quad (3.4.5)$$

$$= \bullet + (4\bullet-)x^2 + (16- \bullet- + 12\bullet--)x^4 + \dots \quad (3.4.6)$$

where the diagrams denote the types of self-retracing walk that contribute at each order.

3.4.2 Open walks

We now generalise this result to include open walks. Consider a walk that begins at site r and ends at site s on \mathcal{B}_z , where r and s are separated by a total of ℓ bonds on the Bethe lattice. We denote the corresponding generating function as $T_{rs}(x)$ [by symmetry, $T_{rs}(x) = T_{sr}(x)$]. We will for convenience draw the Bethe lattice as in Fig. 3.4, the links between r and s (inclusive) forming a backbone, and refer to s as being to the right of r , such that $s = r + \ell$. The walk from $r \rightarrow s$ may then be decomposed in the following way:

- (i) (a) hopping to one of the $z - 1$ nearest neighbours of r not equal to $r + 1$,
- (b) performing a self-retracing walk that begins and ends at depth $k = 1$,
- (c) hopping back to r ,
- (d) performing a walk from r to s .

- (ii) (a) hopping to the ‘right’ of r to site $r + 1$,
 (b) performing a walk from $r + 1$ to s .

The walk from $r + 1$ to s is then decomposed in a similar fashion. Therefore, in terms of the individual generating functions,

$$T_{rs}^{(z)}(x) = \underbrace{(z-1)x^2 R^{(z)}(x) T_{rs}^{(z)}(x)}_{(i)} + \underbrace{x T_{r+1s}^{(z)}(x)}_{(ii)}, \quad (3.4.7)$$

where $R^{(z)}(x)$ is the generating function for self-retracing walks that begin and end at a depth of $k \geq 1$. The labels (i) and (ii) refer to the corresponding steps in the above physical decomposition. From our previous analysis of $T_0^{(z)}(x)$, we know that

$$R^{(z)}(x) = \frac{1}{(z-1)x^2} = \frac{1 - \sqrt{1 - 4(z-1)x^2}}{2(z-1)x^2}, \quad (3.4.8)$$

where again the sign of the square root is chosen to give $R^{(z)}(x) \rightarrow 1$ in the limit $x \rightarrow 0$. The recursion relation (3.4.7) can then be solved to find an expression for $T_{rs}^{(z)}(x)$ in closed form:

$$T_{rs}^{(z)}(x) = \left[x R^{(z)}(x) \right]^\ell T_{ss}^{(z)}(x). \quad (3.4.9)$$

The function $T_{ss}^{(z)}(x)$ which terminates the recurrence relation is simply $T_0^{(z)}(x)$ derived in the previous section, i.e., enumerating the number of perfectly self-retracing paths that begin and end at the same point on the Bethe lattice. We therefore arrive at the final result:

$$T_{rs}^{(z)}(x) = \left(\frac{1 - \sqrt{1 - 4(z-1)x^2}}{2(z-1)x} \right)^\ell \frac{2(z-1)}{z-2 + z\sqrt{1 - 4(z-1)x^2}} \quad (3.4.10)$$

$$\equiv S^{(z)}(x)^\ell T^{(z)}(x). \quad (3.4.11)$$

By virtue of the symmetry of the Bethe lattice, $T_{rs}^{(z)}(x)$ depends only on the length ℓ of the path separating the sites r and s , not on the specific choice of path.

3.4.3 Constrained closed walks

We now further generalise to the case of closed, self-retracing walks on the original lattice \mathcal{L}_z on which the spinon hops in real space. In order to calculate $P_s(t)$, we are required to enumerate the

number of perfectly self-retracing paths that visit the sites $0 \rightarrow s \rightarrow 0$. Any such path can be decomposed as follows:

- (i) a *nonreversing* base path connecting 0 and s on \mathcal{L}_z ,
- (ii) self-retracing excursions which decorate the base path.

The base path must be common to both outward ($0 \rightarrow s$) and return ($s \rightarrow 0$) paths, while the self-retracing excursions can differ between the two paths. In this way, the return path completely “erases” the outwards path, and the path is overall perfectly self-retracing, therefore enclosing precisely zero area. The base paths must be nonreversing, since immediate reversal of the base path corresponds to a self-retracing excursion, which would lead to double counting of such a path. An example of a self-retracing round trip between two sites is shown in Fig. 3.2.

The connection between the Bethe lattice and the original lattice comes from the number of base paths that the particle may take to get between the origin and the site s . Suppose that we know the generating function for the number of nonreversing paths that connect the origin (0), and some other site s on the original lattice \mathcal{L}_z , which we denote by

$$C_s(x) = \sum_{\ell=0}^{\infty} c_{\ell}^{(s)} x^{\ell}. \quad (3.4.12)$$

The generating function for fully self-retracing paths that connect $0 \rightarrow s \rightarrow 0$ can then be constructed in the following way from the three elementary generating functions $S^{(z)}(x)$, $T^{(z)}(x)$ and $C_s(x)$, defined in (3.4.11) and (3.4.12). We will henceforth drop the explicit dependence of these generating functions on the coordination number z for notational convenience. Using generating variables x and y to count the number of steps taken on the outwards and return trips, respectively,

$$\mathcal{P}_s(x, y) = \sum_{\ell=0}^{\infty} c_{\ell}^{(s)} S(x)^{\ell} T(x) S(y)^{\ell} T(y) = T(x) T(y) C_s [S(x) S(y)], \quad (3.4.13)$$

i.e., for each base path, and at each step, a self-retracing excursion may take place, enumerated by the functions $S(x)$ and $T(x)$. As required, the generating function $\mathcal{P}_s(x, y)$ is symmetric under exchange of forwards and backwards walks (i.e., $x \leftrightarrow y$).

3.4.4 Nonreversing walks

Equation (3.4.13) shows that the number of nonreversing base paths on the original lattice, enumerated by $C_s(x)$, is a crucial ingredient in determining the transition probability $P_s(t)$. Our final task therefore is to determine explicitly the generating function $C_s(x)$ (for an arbitrary site s), a general

method for which is presented here. We will introduce the strategy for the square lattice, with the generalisation to the triangular and honeycomb lattices (relating to the quasiparticle excitations on the kagome lattice) deferred to Appendix A.

Since the nonreversing constraint only depends on the previous step in the lattice walk, it may be enforced using $z \times z$ matrices [276]. Let us introduce the generating variables x , δ and ϵ which count the length of the walk, and the number of steps taken in the direction of the (for the square lattice, orthonormal) lattice vectors \mathbf{e}_1 and \mathbf{e}_2 , respectively. At each step, there are four possible directions that the particle may choose from: δ , ϵ , ϵ^{-1} and δ^{-1} . However, for all but the initial step of the walk, the direction which immediately reverses the previous step is forbidden. This may be enforced using the matrix

$$N = x \begin{pmatrix} \delta & \epsilon & \epsilon^{-1} & 0 \\ \delta & \epsilon & 0 & \delta^{-1} \\ \delta & 0 & \epsilon^{-1} & \delta^{-1} \\ 0 & \epsilon & \epsilon^{-1} & \delta^{-1} \end{pmatrix}, \quad (3.4.14)$$

and the initial condition $N_0 = x \text{diag}(\delta, \epsilon, \epsilon^{-1}, \delta^{-1})$. The row index corresponds to the previous step, and the column index to the current step. At each step, the length of the path is advanced by one, and matrix multiplication ensures that all possible combinations of steps are accounted for. The zero entries enforce the nonreversing constraint—any path that immediately reverses its direction is given a coefficient of zero. The initial matrix N_0 imposes that the initial step is unconstrained. Thence the elements of the matrix $N_0 N^{\ell-1}$ give the paths of length ℓ that are consistent with the nonreversing constraint. The full generating function $\mathcal{N}(x; \delta, \epsilon)$ for nonreversing paths is therefore given by the sum over all matrix elements and all possible path lengths ℓ (including also the trivial walk of zero length):

$$\mathcal{N}(x; \delta, \epsilon) = 1 + \sum_{i,j} \sum_{\ell=1}^{\infty} [N_0 N^{\ell-1}]_{ij} = 1 + \sum_{i,j} [N_0 (\mathbb{1}_z - N)^{-1}]_{ij}, \quad (3.4.15)$$

where $\mathbb{1}_z$ is the $z \times z$ identity matrix. Evaluating the inverse of the matrix $\mathbb{1}_z - N$, we arrive at the following explicit expression for the generating function for nonreversing walks on the square lattice

$$\mathcal{N}(x; \delta, \epsilon) = \frac{1 - x^2}{1 + 3x^2 - x(\delta + \delta^{-1} + \epsilon + \epsilon^{-1})}, \quad (3.4.16)$$

consistent with Ref. [276]. This generating function and its counterparts for the other two-dimensional lattices considered in Appendix A represent a central object in this work since they give access to the family of generating functions $C_s(x)$ for all sites s , and hence contain complete information about the spinon density profile after a quench in the magnetic field strength.

Noting that

$$\mathcal{N}(x; \delta, \epsilon) = \sum_{s \in \mathcal{L}_z} \delta^{s_1} \epsilon^{s_2} C_s(x), \quad (3.4.17)$$

the function $C_s(x)$, with $s = \sum_i s_i \mathbf{e}_i$, may be extracted from \mathcal{N} by singling out the terms in (3.4.17) proportional to $\delta^{s_1} \epsilon^{s_2}$. This may be accomplished using the transformation

$$C_s(x) = \int_{-\pi}^{\pi} \frac{d\theta}{2\pi} \int_{-\pi}^{\pi} \frac{d\phi}{2\pi} \mathcal{N}(x; e^{i\theta}, e^{i\phi}) e^{-is_1\theta - is_2\phi}. \quad (3.4.18)$$

Substituting in for the generating function $\mathcal{N}(x; \delta, \epsilon)$, we arrive at the following simplified expression

$$C_s(x) = \frac{(1-x^2)}{2\pi^2 x} \int_0^\pi d\theta \int_0^\pi d\phi \frac{\cos(s_1\theta) \cos(s_2\phi)}{t - \cos\theta - \cos\phi}, \quad (3.4.19)$$

where we have defined $t = (1 + 3x^2)/2x$. This integral may be evaluated by exploiting an equivalence with the Green's function of two dimensional tight-binding models with Hamiltonian \hat{H} . Consider

$$\hat{G}(w) = \sum_{\mathbf{k}} \frac{|\mathbf{k}\rangle \langle \mathbf{k}|}{w - E(\mathbf{k})}, \quad (3.4.20)$$

which satisfies $(w - \hat{H})\hat{G} = 1$. The states $|\mathbf{k}\rangle$ are eigenstates of \hat{H} with energies $E(\mathbf{k})$. Taking matrix elements of $\hat{G}(w)$ with respect to sites $|\mathbf{l}\rangle, |\mathbf{m}\rangle$,

$$G(w; \mathbf{l}, \mathbf{m}) \equiv \langle \mathbf{l} | G(w) | \mathbf{m} \rangle = \frac{1}{\pi^2} \iint_0^\pi \frac{\prod_i dk_i \cos[k_i(l_i - m_i)]}{w - E(\mathbf{k})}. \quad (3.4.21)$$

Hence, when \hat{H} corresponds to a two-dimensional tight-binding model on the square lattice with $E(\mathbf{k}) = \cos k_x + \cos k_y$, we observe the equivalence of (3.4.19) and (3.4.21) up to prefactors, making the identifications $(k_1, k_2) \leftrightarrow (\theta, \phi)$, $w \leftrightarrow t$ and $s_i \leftrightarrow l_i - m_i$.

As shown in, e.g., Refs. [147, 277], the Green's function $G(w; \mathbf{0}, \mathbf{0})$, which is related to the spinon return probability $P_0(t)$, is given exactly by

$$G(w; \mathbf{0}, \mathbf{0}) = \frac{2}{\pi w} K\left(\frac{2}{w}\right), \quad (3.4.22)$$

where K is the complete elliptic integral of the first kind. This result gives rise to the generating function

$$C_0(x) = \frac{2}{\pi} \left(\frac{1-x^2}{1+3x^2} \right) K\left(\frac{4x}{1+3x^2}\right). \quad (3.4.23)$$

The Green's functions for general sites \mathbf{l}, \mathbf{m} (and therefore C_s for a general site s) can also be obtained explicitly using the recursion relations presented in Refs. [147, 277]. This procedure is used later in Sec. 3.5.3 to construct the spatially-resolved spinon density profile.

3.5 Analytical results

Now that we have presented all of the preliminary results, we focus on understanding the high-temperature limits of the physical quantities introduced in Sec. 3.3 that may be inferred from the generating functions for self-retracing walks.

3.5.1 Single spinon density of states

As noted in Sec. 3.3.1, the single-spinon Green's function $G_{ii}(t)$ may at high temperatures be expressed in terms of the exponential generating function $\tilde{T}^{(z)}(x)$ corresponding to closed walks on the Bethe lattice \mathcal{B}_z which, by construction, enclose zero area. The exponential generating function can be constructed from the ordinary generating function $T^{(z)}(x)$ derived in Sec. 3.4.1 using the transformation [see the note on page 93]

$$\tilde{T}^{(z)}(x) = \oint_C \frac{dw}{2\pi i} \frac{e^{xw}}{w} T^{(z)}\left(\frac{1}{w}\right). \quad (3.5.1)$$

The contour C can be shrunk around the branch cut in $(1/w)T^{(z)}(1/w)$ that lies along the real axis between $-2\sqrt{z-1} < \text{Re}(w) < 2\sqrt{z-1}$ *, which gives rise to the expression

$$G_{ii}^{(z)}(t) = \int_{-2\sqrt{z-1}}^{2\sqrt{z-1}} \frac{du}{2\pi} e^{ihtu} \frac{z\sqrt{4(z-1)-u^2}}{z^2-u^2}, \quad (3.5.2)$$

or, equivalently, to the Brinkman and Rice [265] density of states

$$\rho(\omega) = \begin{cases} \frac{z}{2\pi h} \frac{\sqrt{4(z-1)-\omega^2/h^2}}{z^2-\omega^2/h^2} & \text{for } |\omega| < 2\sqrt{z-1}h, \\ 0 & \text{otherwise,} \end{cases} \quad (3.5.3)$$

for single particle excitations.

In the case of the square lattice, for example, this result predicts that the support of $\rho(\omega)$ is narrowed by 13% from $|\omega| < 4h$ at zero temperature to $|\omega| < 2\sqrt{3}h$ at “infinite temperature”, and that $\rho(\omega)$ vanishes like $|2\sqrt{3}h \mp \omega|^{1/2}$ at the positive/negative band edge as opposed to a step-like singularity typical for massive, free quantum particles in two dimensions. The Bethe lattice mapping does not account for the Lifshitz tails, nor any singular behaviour near $\omega = 0$ that has been predicted theoretically [262, 278, 279] and observed numerically [280] in similar models. These differences arise from the neglect of loop diagrams as shown in Fig. 3.3—when a lattice path includes a closed loop, there are two ways in which the loop can be traversed, whereas the Bethe

*The function $T(w^{-1})/w$ does not exhibit a pole at $w = 0$.

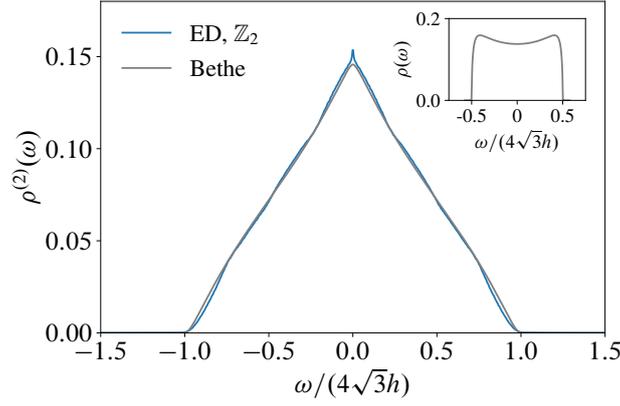


Figure 3.5: Two spinon density of states. The Bethe lattice result corresponds to the autoconvolution of the Brinkman–Rice result (3.5.3), shown in the inset, while the numerical results are calculated using exact diagonalisation (ED) of the effective tight binding Hamiltonian on a system of size $L = 13$ containing two hardcore bosons, satisfying open boundary conditions, averaged over 35 realisations of disorder.

lattice approximation leads to a coefficient of one. Nevertheless, the mapping does capture many of the salient features of the high-temperature density of states over the full spinon bandwidth, as shown in Fig. 3.5. For example, one may use (3.5.3) to understand the behaviour of the density of states with increasing temperature for Majorana fermions in the Kitaev honeycomb model [281].

3.5.2 Single spinon Green’s function

Analogous to the on-site Green’s function, the generating function $T_{ij}(x)$ is related (for $i \neq j$) to the off-diagonal matrix elements of the Green’s function $G_{ij} = \langle \langle \hat{b}_i(t) \hat{b}_j^\dagger(0) \rangle \rangle$ at high temperature. As noted previously, such a quantity is not on its own gauge invariant and must be multiplied by the phases corresponding to a given lattice path γ connecting the sites i and j , i.e.,

$$G_{ij}^{(z)}(t | \gamma) \equiv \left\langle \left\langle e^{i \sum_{\langle \alpha\beta \rangle \in \gamma} \phi_{\alpha\beta}} \hat{b}_i(t) \hat{b}_j^\dagger(0) \right\rangle \right\rangle. \quad (3.5.4)$$

Converting to the corresponding exponential generating function using (3.5.1), the contour can again be shrunk around the branch cut on the real axis between $-2\sqrt{z-1} < \text{Re}(w) < 2\sqrt{z-1}$ and we arrive at the expression

$$G_{ij}^{(z)}(t | \gamma) = [4(z-1)]^{1-\ell/2} \int_0^\pi \frac{d\phi}{2\pi} e^{iht2\sqrt{z-1}\cos\phi} \text{Im} \left\{ \frac{e^{i\ell\phi} \sin\phi}{(z-2)\cos\phi - iz\sin\phi} \right\}, \quad (3.5.5)$$

where ℓ is the length of the path γ from i to j^* .

*Note that the independence of $G_{ij}(t | \gamma)$ on the precise choice of path γ , other than its length ℓ , is a consequence of the self-retracing path approximation.

We note that this quantity also equals the projection of the wave function $|\psi(t)\rangle$ of a particle initially localised at the origin of the Bethe lattice \mathcal{B}_z onto a site at depth ℓ , i.e., $\psi_\ell(t) = \langle \ell | \psi(t) \rangle$, at zero temperature (where the time evolution is generated by a nearest neighbour tight-binding Hamiltonian).

3.5.3 Spinon density profile

We now turn to our main result: characterising the density profile of a spinon initially localised at the origin of the original lattice. The transition probability $P_s(t)$ can be constructed from the generating function $\mathcal{P}_s(x, y)$ in (3.4.13) via conversion to the corresponding exponential generating function using complex contour integration:

$$P_s(t) = \oint \oint \frac{dw_1}{2\pi i} \frac{dw_2}{2\pi i} \frac{e^{iht(w_1-w_2)}}{w_1 w_2} \mathcal{P}_s\left(\frac{1}{w_1}, \frac{1}{w_2}\right), \quad (3.5.6)$$

over sufficiently large circles in both the w_1 and w_2 complex planes, for example. The moments of the probability distribution $P_s(t)$ can then be constructed using the generating functions $T(x)$, $S(x)$ and $C_s(x)$ from Sec. 3.4:

$$\langle \mathbf{r}^{2k}(t) \rangle = \oint \oint \frac{dw_1}{2\pi i} \frac{dw_2}{2\pi i} \frac{e^{iht(w_1-w_2)}}{w_1 w_2} T\left(\frac{1}{w_1}\right) T\left(\frac{1}{w_2}\right) \mathcal{R}_{2k} \left[S\left(\frac{1}{w_1}\right) S\left(\frac{1}{w_2}\right) \right], \quad (3.5.7)$$

where we have defined the function

$$\mathcal{R}_{2k}(x) = \sum_{s \in \mathcal{L}_z} \text{dist}(0, s)^{2k} C_s(x), \quad (3.5.8)$$

with $\text{dist}(0, s) = \sqrt{s_x^2 + s_y^2}$ on the square lattice. Using the relationship (3.4.17) between the generating function for nonreversing walks, $\mathcal{N}(x; \delta, \epsilon)$, and $C_s(x)$, we deduce that the function \mathcal{R}_{2k} may be expressed in terms of appropriate derivatives of \mathcal{N} :

$$\mathcal{R}_{2k}(x) \equiv \sum_{s \in \mathcal{L}_z} (s_x^2 + s_y^2)^k C_s(x) \quad (3.5.9)$$

$$= \left\{ [(\delta \partial_\delta)^2 + (\epsilon \partial_\epsilon)^2]^k \mathcal{N}(x; \delta, \epsilon) \right\} \Big|_{\delta=\epsilon=1}, \quad (3.5.10)$$

which we will write symbolically as $\mathcal{R}_{2k} = [\nabla^{2k} \mathcal{N}](x; 1, 1)$. The expression (3.5.9) and hence (3.5.10) must be generalised to include cross-terms between δ and ϵ if the two basis vectors \mathbf{e}_i are not orthonormal, as is the case for the triangular and honeycomb lattices (see Appendix A for further details).

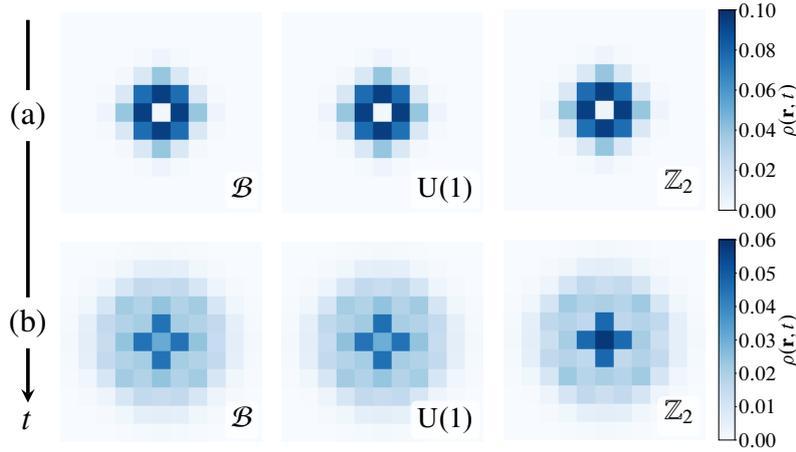


Figure 3.6: At short times, the transition probabilities $P_s(t)$ exhibit coherent oscillatory behaviour. We plot the spinon density profile at times (a) $ht = 0.9$ and (b) $ht = 1.8$, which correspond approximately to a minimum and a maximum of the return probability $P_0(t)$, as predicted by (3.5.6), respectively. From left to right, the profiles correspond to the Bethe lattice analytical result, the continuous flux model, and the π -flux model, labelled \mathcal{B} , $U(1)$, and \mathbb{Z}_2 , respectively. The Bethe lattice mapping provides an essentially exact description of the full density profile for the continuous flux model, and a very good approximation to the π -flux model, at times on the order of the hopping timescale. Notably, discrepancies first become manifest in the π -flux model at the origin due to the stronger localisation. The numerical data for the \mathbb{Z}_2 and $U(1)$ cases are averaged over 25 000 infinite temperature flux configurations.

Short-time dynamics

At sufficiently short times, $ht \lesssim \text{dist}(0, s)$, the transition probabilities $P_s(t)$ exhibit complex oscillatory behaviour arising from the interference of lattice walks of varying lengths, and one must use the full expression (3.5.6) in order to accurately capture the density profile in this regime.

Take for example the return probability $P_0(t)$ obtained using the generating function $C_0(x)$ in (3.4.23). The probability decays with time asymptotically as $P_0(t) \sim 1/t$, suggesting that the spinon asymptotically exhibits diffusive behaviour, which we study in more detail in the next section. However, there also exist superimposed, subleading oscillations due to interference effects that decay as $\cos(4\sqrt{3}t)/t^2$, which may be revealed by applying the method of stationary phase to (3.5.6).

In Fig. 3.6 we plot the spinon density profile predicted by (3.5.6) at $ht = 0.9$ and $ht = 1.8$ [corresponding approximately to extrema of the return probability $P_0(t)$, i.e., $ht \simeq n\pi/(4\sqrt{3})$, for integer n], making use of the family of generating functions $\{C_s(x)\}$, and we compare it with numerical simulations of the disordered tight-binding model (3.2.15) for the case of (i) π -fluxes, and (ii) continuous fluxes, $\phi \in [0, 2\pi)$. We observe almost perfect agreement between the analytical results and the numerics at the shortest of the two times, whereas the quantitative agreement survives at the later time for the continuous flux model only.

Asymptotic second moment

We now focus on the asymptotic behaviour of the density profile, once the transient, oscillatory behaviour of the distribution has subsided. Let us restrict our attention briefly to the second moment of the density distribution, $\langle \mathbf{r}^2(t) \rangle$. For the case of the square lattice, evaluating the derivatives in (3.5.10), one arrives at

$$\mathcal{R}_2(x) = \frac{4x(1+x)}{(1-3x)^2(1-x)}. \quad (3.5.11)$$

Crucially, the function $\mathcal{R}_2(x)$ has a second order pole at $x = (z-1)^{-1}$. This feature is shared by the other lattices considered in Appendix A, and dominates the long-time behaviour of the root mean square (RMS) displacement. In particular, a second order pole in the function \mathcal{R}_2 gives rise to a *linear*, i.e., diffusive, t -dependence of $\langle \mathbf{r}^2(t) \rangle \simeq 2D_z ht$, for sufficiently large times. The full time-dependence of $\langle \mathbf{r}^2(t) \rangle$ described by (3.5.7) corresponds to a crossover from ballistic to diffusive behaviour at a time $ht \sim 1$ (the characteristic time taken for the spinon to hop one lattice spacing). This is because for sufficiently short times the particle has not moved far enough to enclose any flux, and so interference effects do not play a significant role. The linear time dependence at long times is a direct consequence of the result

$$\text{Res}_{w=0} f(w) \frac{e^{itw}}{w^2} = itf(0) + f'(0) \stackrel{t \gg 1}{\approx} itf(0), \quad (3.5.12)$$

if the function $f(w)$ is analytic at $w = 0$. Note that in fact there exists a line of poles along the real axis in (3.5.7) since $S(u+i0^+)S(u-i0^+) = (z-1)$ for $u \in \mathbb{R}$ and $|u| < 2\sqrt{z-1}$. Expanding the integrand for general z about this singular line, we must integrate over the relevant residues between $-2\sqrt{z-1} < u < 2\sqrt{z-1}$ (where the integrand is singular), which defines the function

$$\mathcal{F}(z) \equiv (z-1) \int_{-2\sqrt{z-1}}^{2\sqrt{z-1}} du \frac{4(z-1) - u^2}{z^2 - u^2} \quad (3.5.13)$$

$$= (z-1) \left[4\sqrt{z-1} - z \left(\frac{z-2}{z} \right)^2 \ln \left(\frac{z+2\sqrt{z-1}}{z-2\sqrt{z-1}} \right) \right]. \quad (3.5.14)$$

Comparing the large- t asymptotic expansion of Eq. (3.5.7) with the expected late time behaviour of $\langle \mathbf{r}^2(t) \rangle \sim 2D_z ht$, one therefore obtains the following exact expression for the diffusion constant on a lattice with coordination number z

$$2D_z = \frac{1}{2\pi} \mathcal{F}(z) \lim_{w \rightarrow (z-1)^{-1}} \left(1 - \frac{w}{z-1} \right)^2 \mathcal{R}_2^{(z)}(w). \quad (3.5.15)$$

Lattice	Coordination number, z	D_z
Triangular	6	2.72968
Square	4	2.73383
Honeycomb	3	3.20977

Table 3.1: Values of the diffusion constants D_z obtained by the Bethe lattice mapping corresponding to the long-time behaviour ($ht \gg 1$) of spinons propagating in the high temperature limit ($T > h^4/J^3$).

Using the expression for \mathcal{R}_2 in (3.5.11), we finally deduce that

$$D_4 = \frac{4}{\pi} \left[2\sqrt{3} - \ln(2 + \sqrt{3}) \right] \simeq 2.73383, \quad (3.5.16)$$

for the square lattice. This expression gives the exact value of the diffusion constant observed in, e.g., Ref. [282], which was previously obtained only numerically. The values of the diffusion constant for the triangular and honeycomb lattices are given in Tab. 3.1.

In fact, the following expression for D_z is valid for all three lattices

$$\frac{D_z}{\langle d^2 \rangle} = \frac{z^2}{4\pi(z-2)} \left[4\sqrt{z-1} - z \left(\frac{z-2}{z} \right)^2 \ln \left(\frac{z + 2\sqrt{z-1}}{z - 2\sqrt{z-1}} \right) \right], \quad (3.5.17)$$

normalised by the arithmetic mean of the squared distances, $\langle d^2 \rangle$, corresponding to the possible moves at each step*. A plot of this function in Fig. 3.7 shows that D_z exhibits a minimum at $z \simeq 4.833$ —this is due to the competition between (i) reduced destructive interference at low coordination numbers (vanishing destructive interference as $z \rightarrow 2^+$, since there are no loops for $z = 2$), and (ii) a greater number of paths between any two points on the lattice for large z , the effect of which dominates at large z . This result shows that it is a fortuitous coincidence that the square and triangular lattices exhibit such similar diffusion constants.

We stress that (3.5.17) corresponds to *quantum* diffusion, which results from the complex interference pattern produced by the multitude of lattice walks, and is faster than the corresponding classical diffusion† (random walk), $D_{\text{cl}} = z/2$, for coordination number $z \leq 5$. For $z > 5$, interference effects dominate, leading to slower propagation. The difference between classical and quantum diffusion is further reflected in the non-Gaussian nature of the asymptotic density profile. The calculations presented in this section can be extended to calculate arbitrary moments of the spinon density profile, as shown in Appendix B, from which one can demonstrate explicitly that

*This normalisation allows us to apply Eq. (3.5.17) to lattices such as the double triangular lattice ($z = 8$), which include steps of different lengths.

†We compare our quantum diffusive behaviour with the classical random walk in which the particle hops to one of its neighbours in a time $(zh)^{-1}$. The operators $b_i(t)$ and the classical probabilities then obey the same equations of motion (the latter in imaginary time). This normalisation is consistent with that used in Ref. [282].

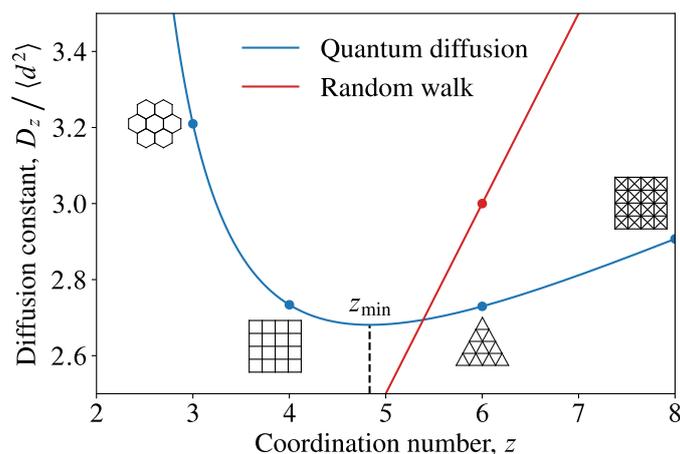


Figure 3.7: A plot of the diffusion constant D_z against the coordination number z from (3.5.17), having set the nearest neighbour distance between lattice sites equal to unity. The corresponding classical diffusion constant $D_z^{\text{cl}} = z/2$ is also shown for comparison. The markers denote the values of D_z for some common lattices: the honeycomb, square, triangular, and double triangular lattices, from left to right.

higher order cumulants do not vanish.

3.6 Numerical results

The Bethe lattice result is compared to Trotterised time evolution [271] generated by the single-particle Hamiltonian (3.2.15) on a 1999×1999 square lattice. The method of Trotterised time evolution was chosen for its ability to simulate system sizes significantly larger than those accessible to exact diagonalisation. The results, which are shown in Fig. 3.8, are averaged over infinite temperature disorder realisations, both for random discrete fluxes ($2\pi/n$, for $n = 2, 3, 4$), as well as for the continuous random flux model, where the flux threading each plaquette is chosen from a uniform distribution $\phi \in [0, 2\pi)$. We see that, over numerically accessible, intermediate timescales, the Bethe lattice result provides an excellent quantitative description of the density profile for the continuous flux model, and qualitative agreement with the \mathbb{Z}_2 model. The results for discrete fluxes with $n > 2$ can be seen to rapidly converge to the continuous flux result. Note that the case $n = 2$, corresponding to \mathbb{Z}_2 fluxes, is special, being the only case in which the effective tight-binding Hamiltonian exhibits time reversal symmetry (i.e., $H = H^*$).

In the π -flux model, encircling a flux an even number of times gives rise to constructive interference. The difference in behaviour between this model and the continuum case means that the loop diagrams depicted in Fig. 3.3b, which are missed by the Bethe lattice, play an important role. These diagrams lead to an increased weight near the origin, which results in a reduced RMS displacement. These observations are consistent with the idea that all single-particle eigenstates of

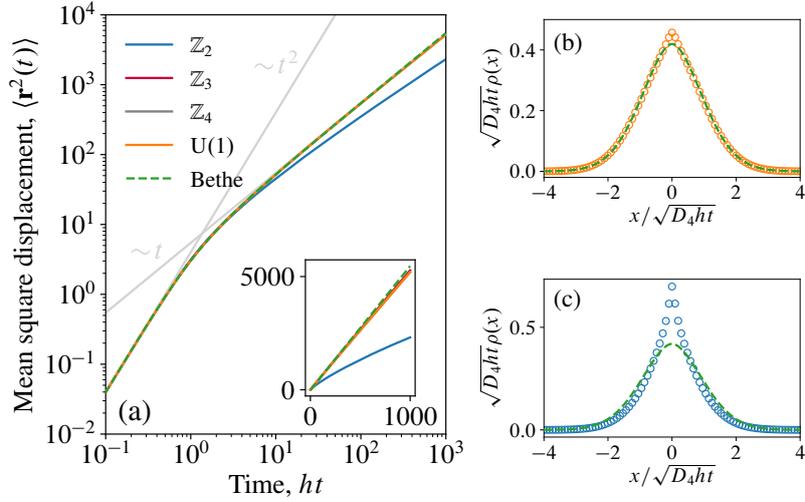


Figure 3.8: (a) Comparison of Bethe lattice result for the mean square deviation, $\langle \mathbf{r}^2(t) \rangle \simeq 2D_z ht$, with numerics for (i) the discrete flux model with fluxes $2\pi/n$, for $n = 2, 3, 4$, and (ii) the continuous random flux model. Time evolution is performed using a high order Suzuki–Trotter decomposition on a square lattice with 1999×1999 sites and averaged over 128 random flux configurations at infinite temperature. All models exhibit the same ballistic behaviour for sufficiently short times. The lines $4t^2$ and $2D_z t$ are shown in light grey as a guide to the eye to indicate the ballistic and diffusive regimes. The π -flux model shows the most pronounced deviation from the Bethe lattice approximation at longer times as a result of being most strongly localised. The density profile at $ht = 40$ for the Bethe lattice is compared with the continuous and π -flux models in (b) and (c), respectively. The profiles $\rho(x)$ are averaged over 10^5 flux configurations on a lattice of size 249×249 sites. The error bars are in all cases too small to be visible.

the effective disordered tight-binding model are localised, but with a diverging localisation length near $E = 0$ arising from the presence of purely off-diagonal disorder in the Hamiltonian.

The subdiffusive form of the second moment observed in the numerics may be explained in the following way: a wave packet composed of states in the vicinity of energy E will diffuse with diffusion constant $D(E)$ up to (approximately) their localisation length $\lambda(E)$ [283], at which point such states give rise to a fixed, time-independent contribution to $\langle \mathbf{r}^2(t) \rangle \sim \lambda(E)^2$. As time progresses, a reduced fraction of states have not yet reached their localisation length and are still diffusing, explaining the negative curvature observed in Fig. 3.8. The Bethe lattice result can therefore be thought of as giving the behaviour of $\langle \mathbf{r}^2(t) \rangle$ before any of the states have reached their localisation length. Since the localisation length of the continuous flux model increases exponentially with energy away from the band edge [280], we observe very close agreement up to $ht = 10^3$ with the Bethe lattice result. The difference in behaviour between the continuous and π -flux models may be attributed to the fact that the latter is more strongly localised [284], implying a smaller fraction of diffusing states at any given time, and hence a more pronounced departure from the pure diffusion predicted by the Bethe lattice.

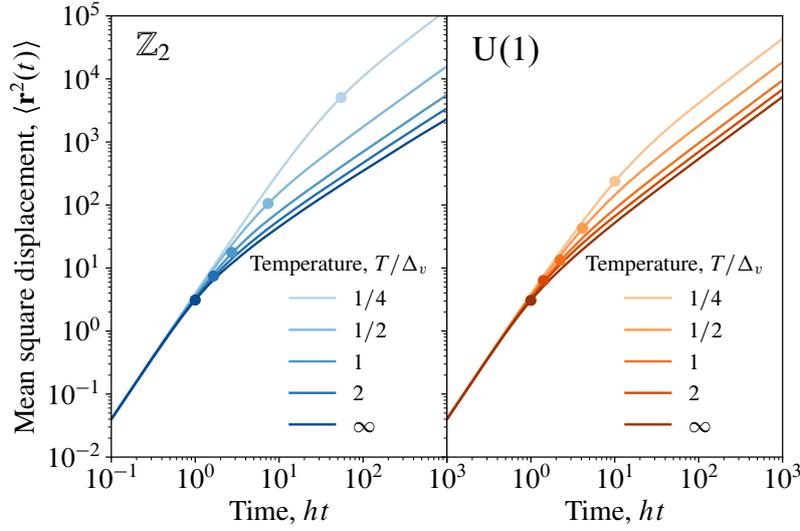


Figure 3.9: Comparison of the mean square displacement $\langle \mathbf{r}^2(t) \rangle$ of a spinon initially localised at the origin for various temperatures. Over the simulated timescales, the dominant effect of reducing temperature is to shift the crossover from ballistic to (sub-) diffusive behaviour to larger times, namely at a time $ht \sim \xi^2$ indicated by the circular markers as a guide to the eye. In the \mathbb{Z}_2 case, ξ is set by the distance between visons. In the $U(1)$ case, ξ^2 is given by the area whose enclosed flux has an $O(1)$ variance, which we regularise at large temperatures using $\xi_c^4 = 1 + 4\beta^2 \Delta_v^2$. Time evolution is performed using a high order Suzuki–Trotter decomposition on a 1999×1999 square lattice, and the data are averaged over 128 disorder realisations for each temperature. The statistical error in the data is smaller than the line width.

3.6.1 Finite temperature

In the case of discrete π -fluxes with a gap $\Delta_v \sim h^4/J^3$, at intermediate temperatures or, equivalently, finite vison separation $\xi \sim \rho^{-1/2} \sim e^{\Delta_v/2T}$, the crossover from ballistic to subdiffusive behavior is shifted to later times. We expect that the particle should propagate ballistically until it has encountered a sufficient number of visons so as to impede its motion: $(ht)^2/\xi^2 \sim ht$. Hence, for $ht \ll \xi^2$ we expect to see free-particle behaviour, and for $ht \gg \xi^2$ we expect to observe approximately the infinite temperature (sub-) diffusive behaviour (with a renormalised diffusion constant). The crossover between the two regimes is therefore set by the time taken to diffuse to the nearest vison. This behaviour is indeed seen in Fig. 3.9.

In the continuous flux model, one may attribute an energy cost $E(\phi) = -\Delta_v \cos \phi$ to threading a given plaquette with a flux ϕ . At sufficiently low temperatures, $\beta\Delta_v \gg 1$, the corresponding probability density $p(\phi) \propto e^{-\beta E(\phi)}$ is approximately Gaussian, and the relevant length scale ξ_c becomes $\xi_c^2(T) = 2/T$ [cf. Eq. (3.3.4)]. This characteristic area is defined via the relation $\langle e^{i \sum_{(\alpha\beta) \in \gamma} \phi_{\alpha\beta}} \rangle \equiv e^{-A_\gamma/\xi_c^2}$, and may be understood as the area such that typical fluctuations of the enclosed flux have a magnitude that is $O(1)$. As in the discrete flux case, the effect of finite temperature is to shift the crossover from ballistic to (sub-)diffusive behaviour to a time $ht \sim \xi_c^2$, as

shown in Fig. 3.9.

3.7 Conclusions

In this chapter, we studied the effects of nontrivial mutual statistics on the propagation of quasiparticles in gapped topological systems at finite temperature. Specifically, we considered a temperature regime where one species of quasiparticle is thermally excited and provides a static (à la Born–Oppenheimer) stochastic background for the other species, which are sparse and hop coherently across the lattice. This is a regime of experimental interest in topological quantum spin liquids, where a large separation of energy scales between different species of quasiparticle arises naturally in many realistic model Hamiltonians.

We used a combination of numerical and analytical approaches to investigate toric-code-inspired toy models, where the excitations (dubbed spinons and visons) have anyonic mutual statistics. The effect of nonzero temperature in our model is to populate a finite density of static visons. Due to the mutual statistics of the quasiparticles, visons act as Aharonov–Bohm half flux quanta for the spinons. Within perturbation theory, our model permits an effective description in which the spinons evolve in time according to a two-dimensional tight-binding Hamiltonian in the presence of randomly placed fluxes. Changing temperature alters the density of the fluxes, which, in turn, changes the strength of off-diagonal disorder in the tight-binding Hamiltonian. We also considered models in which the flux threading each plaquette is a multiple of $1/3$ or $1/4$ of the flux quantum, and the case in which the flux is distributed continuously.

Various time-dependent observables for lattice systems, including the spinon density profile in our effective tight-binding description, may be computed by counting discrete lattice paths. In order to make analytical progress, we considered the self-retracing path approximation. Such paths are expected to dominate at intermediate times due to interference effects by virtue of the Aharonov–Bohm effect. To this end, we map the self-retracing paths to walks on an auxiliary Bethe lattice and enumerate such walks exactly. This gives us access to analytical expressions for the spinon density profile as a function of space and time.

For sufficiently short times, namely on the order of the hopping timescale, τ , the density exhibits oscillatory behaviour due to coherent interference effects. On these timescales, the self-retracing path approximation is essentially exact and our results are almost indistinguishable from numerical simulations. At times much greater than the hopping timescale, $t \gg \tau$, the self-retracing path approximation predicts asymptotic quantum diffusive behaviour of the spinon, i.e., $\langle \mathbf{r}^2(t) \rangle \simeq 2D_z t / \tau$. We obtained an exact expression for the corresponding diffusion constant D_z , which depends on the coordination number of the underlying lattice. The function D_z exhibits a minimum at $z \simeq 5$, where the effects of (i) reduced destructive interference at low coordination

numbers, and (ii) an increasing number of paths connecting any two sites at larger coordination numbers, balance one another. The higher moments of the density distribution in the large-time limit exhibit non-Gaussian behaviour, which highlights the difference between quantum and classical diffusion.

Comparison with numerical simulations reveals excellent agreement with the continuous flux model up to $O(10^3)$ hopping times, while for the π -flux model discrepancies become apparent at much shorter times. This difference is understood as arising from the distinct localisation properties of the two models. When considering $\langle \mathbf{r}^2(t) \rangle$, states with a given energy will diffuse with some characteristic diffusion constant until the corresponding localisation length is reached. The π -flux model is more strongly localised and so at any given time a larger fraction of states have reached their localisation length, and give rise thereafter to a time-independent contribution to $\langle \mathbf{r}^2(t) \rangle$.

The results that we have presented provide us with a quantitative understanding of the crossover from ballistic to quantum (sub-) diffusive motion of spinons through a sea of thermally-excited visons, which is a direct consequence of their nontrivial mutual statistics. More generally, our work represents a step forward in understanding the dynamics of quantum spin liquids at finite temperature, which is essential to interpret both the relevant experiments and numerical data. Our results demonstrate another way in which the mutual semionic statistics of spinons and visons manifests itself in the dynamical properties of spinons; this paves the way for the possible study of such dynamics as an experimentally viable diagnostic tool for anyonic statistics in many-body systems that exhibit topological order.

We expect that our results may be relevant to several interlaced but distinct contexts of many-body physics. On the condensed matter physics front, while realistic Hamiltonians require including further effects, such as possible interactions between quasiparticles and correlations in the spin background, it is nonetheless tempting to point at the recent experimental advances in the study of Kitaev-model-like candidate materials at finite temperature as a possible context where the physics discussed in our work may be relevant and observable [285]. However, to make such connections, some modification of our present framework is necessary in that the spinon dispersion in the Kitaev model [46] is massless and relativistic, as opposed to the massive and non-relativistic dispersion considered in the present work.

In the context of quantum information and quantum computing, the recent proposal that the toric code and similar \mathbb{Z}_2 spin liquid Hamiltonians may be realised using quantum annealers [286], indeed in the limit explored in our work of a large star constraint and a perturbative transverse field, promises to provide further avenues to benchmark and explore the type of phenomena that we have uncovered, in a convenient and highly tunable setting.

Seen from the ultracold atomic physics perspective, our results also describe quantitatively the motion of holes in real space in the large- U , large- S limit of the Hubbard model [265, 282, 287, 288]. As the hole moves throughout the spin environment, it permutes the spins. This “interaction” with

the spin environment leads to dissipationless decoherence [289, 290]—the propagation of the hole is significantly slowed despite there being no transfer of energy between the hole and the spins. Owing to recent developments in quantum gas microscopy [291–296], our analytical expressions for the site-resolved density profile may thus be probed directly in ultracold atomic experiments. Our calculation extends the self-retracing path approximation used in Ref. [287], showing that it in fact holds to much larger times and distances.

Our work also points at a couple of potentially interesting future directions. Extending the analysis in the present work to string-net models [297] may offer access to richer varieties of topological order and anyonic statistics, and an opportunity to classify more generally the resulting dynamics. Perhaps more interesting is the spinons’ back action on the visons, which we have so far neglected. In a similar spirit to Ref. [282], it is plausible that the quantum coherent hopping of spinons may lead to nontrivial correlations in the positions of the visons, which may have other important implications at finite temperature, indicative of quantum spin liquid behaviour. Indeed, in Ref. [298], I showed how spinons can become localised on temperature-dependent patches of expelled visons, in a manner analogous to the Nagaoka effect [282, 299, 300], with significant consequences for the thermodynamic and transport properties of the system.

4

Thermal quenches in classical spin ice

4.1 Motivation

In previous chapters, we have focused on models of quantum spin liquids, where nonzero temperatures lead to the thermal excitation of point-like quasiparticles with exotic properties. These defects were shown to drastically alter the properties of the models in question with respect to their zero temperature behaviour, both in equilibrium and dynamically. Here, we change tack and address the nonequilibrium dynamics of *classical* spin ice materials [301]. The nonequilibrium behaviour of these systems is particularly rich and exciting and they can exhibit remarkably long relaxation and response time scales at low temperatures. While a number of attempts have been made to model and understand the origin of the dynamical behaviour in spin ice materials, the complete picture arguably remains beyond our grasp.

In this chapter, we make progress by investigating the specific setting of thermal quenches in classical spin ice [302], where these systems have been shown to enter long-lived metastable states in which the monopole excitations form so-called noncontractible pairs* (see Fig. 4.3). While the nature of these states is well understood, the dynamical mechanisms underpinning their formation remain hitherto unexplored and are the subject of this work. Using a combination of numerical simulations and analytical mean field theory, we are able to provide a complete understanding of the phenomenon. We find that the emergence of the plateau is rooted in two key ingredients: (i) the long-range nature of the Coulomb interaction between the monopoles; and (ii) the fact that low temperature thermal quenches in spin ice can give rise to a non-hydrodynamic regime that increases the decay rate of the free monopole density in the system. The latter feature is notably reminiscent of jamming and some instances of glassiness. A change in the short-time dynamics of the system allows it to enter a metastable state, which would have been otherwise avoided and whose lifetime can easily become exceptionally long at (experimentally relevant) low temperatures.

Our results further demonstrate that the plateau reduces to a finite size effect when considering a

*A noncontractible pair corresponds to a pair of oppositely charged monopoles residing on adjacent tetrahedra which are unable to annihilate by flipping the intervening spin.

model with solely finite-range interactions between monopoles. Hence, the experimental observation of a metastable plateau corresponding to a finite density of noncontractible pairs in spin ice is direct evidence of the long-range nature of the interactions between the monopoles. This adds one important experimental avenue to study these interactions, whose range has thus far been probed only via the field-dependence of unbinding of monopole pairs [303], and indirectly via the appearance of a liquid-gas phase diagram [88].

Our findings are particularly timely thanks to the recent experimental claim that a state rich in noncontractible pairs can be generated in classical spin ice materials $\text{Dy}_2\text{Ti}_2\text{O}_7$ and $\text{Ho}_2\text{Ti}_2\text{O}_7$ [303] using a so-called avalanche quench protocol [304].

The chapter is structured as follows. We begin by providing a broad introduction to Monte Carlo methods, how they may be used to describe quantitatively the dynamics of classical spin ice, and how continuous time Monte Carlo algorithms provide access to long times at low temperatures when the the system's dynamics slow down. We then review the background on thermal quenches in classical spin ice and summarise the main results obtained in this chapter in Sec. 4.3. We provide an overview of the models we consider in Sec. 4.4, and present our Monte Carlo results in Sec. 4.5, including a finite size scaling analysis of the density of noncontractible pairs in the metastable plateau. Section 4.6 is devoted to the use of mean field population dynamics to understand the differences in behaviour between the various models and types of interaction. We draw our conclusions and highlight the relevance of our results to experiments in Sec. 4.7.

4.2 Monte Carlo methods

Suppose that we are interested in describing the (for now) equilibrium properties of a classical system with Hamiltonian H . In a non-driven setting, the Hamiltonian simply assigns an energy to each state of the system. Let us assume that our system is described by a discrete set of states, labelled by the index μ , with energies E_μ . The dynamics of the system as it makes transitions between the various energy levels is described by the master equation

$$\frac{dw_\mu}{dt} = \sum_\nu [w_\nu(t)R_{\nu\rightarrow\mu} - w_\mu(t)R_{\mu\rightarrow\nu}] , \quad (4.2.1)$$

where $w_\mu(t)$ is the probability that the system occupies the state μ at time t [305, 306]. The real and non-negative quantities $R_{\mu\rightarrow\nu}$ describe the rates at which the system makes transitions between the different discrete states. Normalisation of the probability distribution implies that $\sum_\mu w_\mu(t) = 1$ at all times. When the system is in equilibrium, the left hand side of (4.2.1) vanishes, and the probability distribution w_μ is stationary in time. For a system in contact with a thermal reservoir at

temperature $T = \beta^{-1}$, the stationary distribution is the Boltzmann distribution*

$$\lim_{t \rightarrow \infty} w_\mu(t) = \frac{1}{Z} \exp(-\beta E_\mu), \quad (4.2.2)$$

where $Z = \sum_\mu e^{-\beta E_\mu}$ is the partition function. The equilibrium expectation value of some quantity A , denoted $\langle A \rangle$, is equal to

$$\langle A \rangle = \frac{1}{Z} \sum_\mu A_\mu e^{-\beta E_\mu}, \quad (4.2.3)$$

where the quantity A assumes the value A_μ in the state μ . Physically, the expectation value can be interpreted either as (i) an ensemble average over many statistically independent copies of the system, or (ii) a time average of $A(t)$ in a single copy of the system (as one would have in an experimental setting). The equivalence of these two interpretations only holds when the time average is taken over a sufficiently long period of time so as to sample many “independent” configurations of the system. If, conversely, the time average is taken over a period of time that is too short, the system may not have explored a representative sample of the states that are available at the temperature of the system [305].

The above recipe, provided to us by statistical mechanics, tells us how to compute the average value of an observable A in thermodynamic equilibrium. However, the total number of states that the system may occupy is typically intractably large. Take, for example, the particularly simple example of a d -dimensional classical Ising model. The spin on each lattice site may point up or down, implying that the total number of states scales as $D = 2^{L^d}$ in a system with linear dimension L . In a three dimensional system with just $L = 3$ this number is already $D = 134\,217\,728$. In the numerical evaluation of expressions like (4.2.3), it is therefore numerically infeasible to perform the summation by exact enumeration of all states for all but the smallest system sizes. Instead, we would like to approximate $\langle A \rangle$ by taking a finite number of states, say $M \ll D$, from a uniform probability distribution over μ . In this case,

$$\langle A \rangle \stackrel{?}{\approx} \frac{\sum_{\ell=1}^M A_{\mu_\ell} e^{-\beta E_{\mu_\ell}}}{\sum_{\ell=1}^M e^{-\beta E_{\mu_\ell}}}. \quad (4.2.4)$$

However, this approximation is, in general, rather insufficient to capture the behaviour of the system. This becomes particularly apparent at low temperature (large β), where only the lowest energy states are sampled. Such states correspond to an exponentially small region of state space, and so the uniform sampling in (4.2.4) correspondingly requires an exponentially large M in order to sample the atypical, low energy states. If, instead, we are able to sample directly from the Boltzmann

*Indeed, it is possible to show that the quantum master equation from Sec. 1.3.1 maps onto an equation of the same form as (4.2.1), where the probabilities w_μ correspond to the diagonal elements of the system’s density matrix in the eigenbasis of the system’s Hamiltonian [23].

distribution (4.2.2), the expectation value reduces to

$$\langle A \rangle \approx \frac{1}{M} \sum_{\ell=1}^M A_{\mu_\ell}. \quad (4.2.5)$$

This expression corresponds to a particular example of *importance sampling*. Unlike Eq. (4.2.4), a prohibitively large number of states M are no longer required since we are most likely to pick the states in which the system spends most of its time at a given temperature.

In order to sample directly from the Boltzmann distribution, we make use of *Markov processes*. A Markov process consists of a set of time independent transition probabilities $P_{\mu \rightarrow \nu}$ from state μ to states ν . These probabilities depend only on the current state of the system, not on its history (the system has no memory of the states that it has passed through to get to the state μ). In order to reproduce the Boltzmann distribution, the Markov process must be *ergodic*. This means that no state of the system should be inaccessible. Equivalently, the system will eventually visit every state, given sufficient time (although, in practice, this time may be astronomically large). Conservation of probability requires that the rate of transitions out of the state μ must equal the rate of transitions back into it:

$$p_\mu = \sum_{\nu} p_\nu P_{\nu \rightarrow \mu}. \quad (4.2.6)$$

However, this condition alone is insufficient to *guarantee* that the Boltzmann distribution will be the only equilibrium state for the Markov process. To mitigate the possibility of obtaining limit cycles, one applies a more stringent constraint on the dynamics than merely conservation of probability. This condition is known as *detailed balance*, which requires that the transition probabilities satisfy [305–307]

$$p_\mu P_{\mu \rightarrow \nu} = p_\nu P_{\nu \rightarrow \mu}. \quad (4.2.7)$$

Although the above condition is now sufficient to reach equilibrium, it is not strictly necessary, and in some cases equilibrium can be reached more rapidly by violating detailed balance (see, e.g., Refs. [308, 309] and references therein). Often, one further decomposes the transition probability as $P_{\mu \rightarrow \nu} = g_{\mu \rightarrow \nu} A_{\mu \rightarrow \nu}$, into a selection probability, $g_{\mu \rightarrow \nu} \leq 1$, and an acceptance ratio, $A_{\mu \rightarrow \nu} \leq 1$. As long as ergodicity is satisfied, the transition probabilities can be chosen in any way that one desires, while the acceptance ratios then ensure that detailed balance is satisfied

$$e^{-\beta(E_\mu - E_\nu)} = \frac{P_{\nu \rightarrow \mu}}{P_{\mu \rightarrow \nu}} = \frac{g_{\nu \rightarrow \mu} A_{\nu \rightarrow \mu}}{g_{\mu \rightarrow \nu} A_{\mu \rightarrow \nu}}. \quad (4.2.8)$$

Expectation values of the form (4.2.3) may then be evaluated by following the time evolution of the Markov process and (once it has reached equilibrium) taking “measurements” separated by appropriate time intervals.

As well as being an indispensable numerical tool for calculating the properties of systems in thermal equilibrium, the Monte Carlo methods described above can also be used to simulate the finite temperature dynamics of classical systems. Indeed, in order to sample from the Boltzmann distribution, we have constructed a discrete-time version of the master equation (4.2.1) with the same stationary distribution. If the proposed updates, specified by the $g_{\mu \rightarrow \nu}$, are in some sense local, then the evolution of the Markov process can provide a good description of the true equilibrium dynamics of the system (once one finds an appropriate way of converting between real time and Monte Carlo time). This is the case in spin ice, whose dynamics is well approximated at low temperatures ($T \lesssim 5$ K) by incoherent spin reversals [92], which is mimicked in Monte Carlo simulations by single spin flip dynamics.

Nonlocal cluster updates, whilst beneficial in the context of speeding up the approach to equilibrium, are not appropriate for describing the real-time evolution of classical spin systems. In the setting of the two-dimensional Ising model, for example, the critical slowing down [310] that occurs in the vicinity of the critical point, which cluster updates (e.g., the Wolff algorithm [311]) are able to efficiently bypass, is a physical divergence of the time scale relevant to the system's dynamics.

4.2.1 Single spin flip dynamics

In classical spin ice, the relevant degrees of freedom are classical Ising spins $S_i \in \{-1, +1\}$, as explained in Sec. 1.4.3. In a single spin flip update, there are N_s possible spin flips, and the selection probability $g_{\mu \rightarrow \nu} = N_s^{-1}$, independent of the initial state μ and the final states ν . Detailed balance is encoded in the acceptance ratio, $A_{\mu \rightarrow \nu}$. The Metropolis algorithm corresponds to the choice

$$A_{\mu \rightarrow \nu} = \begin{cases} e^{-\beta(E_\nu - E_\mu)} & \text{if } E_\nu - E_\mu > 0 \\ 1 & \text{otherwise.} \end{cases} \quad (4.2.9)$$

That is, we always accept the move if it lowers the energy of the system, and if the move increases the energy of the system by $E_{\mu\nu} = E_\nu - E_\mu > 0$, then the move is accepted with probability $e^{-\beta E_{\mu\nu}}$. One Monte Carlo *sweep* of the system equals N_s independent single spin flip updates. The prescription for converting between real time and Monte Carlo time is determined by AC susceptibility measurements [312], which suggest that one Monte Carlo sweep corresponds to approximately 1 ms.

4.2.2 Waiting time Monte Carlo

The dynamics of classical systems at low temperatures, especially those that exhibit glassy behaviour, can become very slow. In our simulations, slow dynamics arises in part due to the formation of

metastable noncontractible pairs of monopoles. This manifests as an acceptance rate that becomes undesirably low from a practical point of view. As an extreme example, consider an Ising model $H_{\text{Ising}} = -J \sum_{\langle i,j \rangle} S_i S_j$ at a temperature $T \ll J$ in its ferromagnetic ground state. On a lattice with coordination number z , a spin flip costs an energy $2zJ \gg T$, and thence the acceptance ratio $A_{0 \rightarrow 1} = e^{-2\beta zJ} \ll 1$, where the subscripts correspond to the number of flipped spins relative to fully magnetised background. One therefore has to reject an exponentially large number of proposed updates in order to change the state of the system. For a three-dimensional cubic Ising model with $T/J = 2$, $A_{0 \rightarrow 1}^{-1} \approx 163\,000$. Whilst this slowing down is eminently physical, it presents a challenge to the numerical simulation of such systems.

One way to circumvent this “problem” is to use a rejectionless algorithm. In the above example, rather than attempting $O(10^5)$ spin flips, we can simply perform the spin flip, and update the “time” of the simulation by the length of time it would have taken were we to use the Metropolis algorithm. This is the idea behind the waiting time method [313, 314] (and, more generally, continuous time Monte Carlo algorithms [315]).

Here we summarise the necessary details of the waiting time method, as outlined in Ref. [313]. In the waiting time method, for each spin S_i there exists an associated time variable t_i , which equals the time at which the spin would flip, given its instantaneous local field. The global time is denoted t_g . The flipping times are initially drawn from an exponential distribution

$$P(t_i) = \frac{1}{\tau_i} e^{-t_i/\tau_i}, \quad (4.2.10)$$

with mean $\tau_i = \max(1, e^{\beta \Delta_i})$, where Δ_i equals the energy required to flip the spin S_i . Once the system has been initialised, one repeats the following steps:

1. flip the spin S_j whose flipping time t_j is the smallest: $t_j = \min_i t_i$,
2. update the global time of the simulation according to $t_g = t_j$,
3. generate new waiting times δ_i for the flipped spin and those that interact with it, and update their flipping times $t_i \rightarrow t_g + \delta_i$.

The above algorithm satisfies detailed balance, and also exhibits a dynamical equivalence with the Metropolis algorithm. In the waiting time algorithm, the probability that the system makes a transition from state μ to ν , which differ by a single spin flip at site j , is equivalent to the probability that t_j is the smallest of all $\{t_i\}$. Therefore,

$$p_j(\mu) = \int_0^\infty dt_j \left(\prod_{k \neq j} \int_{t_j}^\infty dt_k \right) \frac{e^{-\sum_i t_i/\tau_i(\mu)}}{\prod_i \tau_i(\mu)} = \frac{\tau_j^{-1}(\mu)}{\sum_i \tau_i^{-1}(\mu)}. \quad (4.2.11)$$

For the Metropolis algorithm, the probability that the system transitions from state μ to state ν , given that a spin flip has occurred, is

$$p_j(\mu) = \frac{P_{\mu \rightarrow \nu}}{\sum_{\nu \neq \mu} P_{\mu \rightarrow \nu}}. \quad (4.2.12)$$

Since $P_{\mu \rightarrow \nu} = \min(1, e^{-\beta(E_\nu - E_\mu)})/N_s = \tau_j^{-1}(\mu)/N_s$ for $\mu \neq \nu$, we observe that (4.2.11) and (4.2.12) are equivalent. The quantities $p_j(\mu)$ and $P_{\mu \rightarrow \nu}$ differ in the exclusion (inclusion) of $P_{\mu \rightarrow \mu}$ in the denominator. The average time taken for a spin flip in the waiting time method is

$$\langle \min_i t_i \rangle = \left(\sum_i \tau_i^{-1}(\mu) \right)^{-1}. \quad (4.2.13)$$

This result is to be contrasted with the survival probability with Metropolis dynamics after a time t (that is, t sweeps of the system, or $n = tN_s$ single spin flip updates)

$$P_{\mu \rightarrow \mu}^n = e^{tN_s \log P_{\mu \rightarrow \mu}}. \quad (4.2.14)$$

The characteristic time scale of the exponential decay is

$$\tau(\mu) = -\frac{1}{N_s \log P_{\mu \rightarrow \mu}} \simeq \frac{1}{N_s \sum_{\nu \neq \mu} P_{\mu \rightarrow \nu}}, \quad (4.2.15)$$

where the second equality is approximate, and holds when $\sum_{\nu \neq \mu} P_{\mu \rightarrow \nu} \ll 1$. That is, the two timescales (4.2.13) and (4.2.15) coincide when the system remains “stuck”, and $\tau(\mu) \gg N_s^{-1}$.

One disadvantage of the waiting time method is the requirement to know the waiting time for *all* states that are connected to the current state of the system at any given time. However, this increase in computational complexity is outweighed by the aforementioned benefits at sufficiently low temperatures. When the interactions between the constituent spins are strictly short ranged, the number of waiting times that need to be generated in the third step does not scale with system size, and the computational complexity per step is significantly reduced.

4.3 Background and summary of results

Dipolar spin ice systems have been predicted to exhibit dynamically-arrested, monopole-rich, metastable states following appropriate thermal and field quenches [302, 316]. Reference [302] recognised that at the heart of the dynamical arrest lies the formation of so-called noncontractible pairs: a monopole and an antimonopole sitting next to one another, separated by a spin whose reversal does not lead to their annihilation. As a result, the two defects become bound to one another

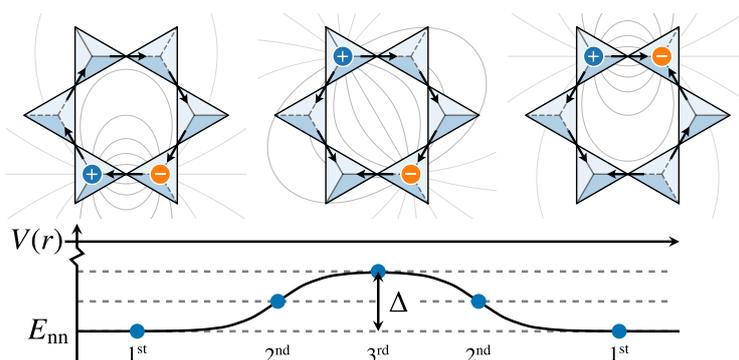


Figure 4.1: Schematic depiction of a noncontractible monopole–antimonopole pair, responsible for the metastable plateau in monopole density observed following a thermal quench to low temperatures in classical spin ice. The activated decay of the pair requires separating its members up to third-neighbour distance, as shown in the central figure, costing an energy Δ (in isolation) due to their mutual Coulombic attraction. The pair is then able to annihilate elsewhere on the lattice, as shown for example in the rightmost figure.

and are unable to move throughout the lattice without separating—a process that costs Coulomb energy due to the mutual attraction between the two opposite charges*. This activation energy barrier explains why a noncontractible pair *per se* is metastable.

In general, two decay channels are available to noncontractible pairs. Firstly, they can separate and annihilate somewhere else on the lattice at the cost of paying an activation energy barrier; the smallest barrier associated with such activated decay processes requires separating the pair up to third-neighbour distance, as shown in Fig. 4.1. Alternatively, pairs can undergo monopole-assisted decay: When the pair is hit by a stray (free) monopole, this causes the annihilation of the oppositely charged member of the pair, thus freeing up its partner [302], as in Fig. 4.2. This second process does not incur an energy barrier and does not change the density of free monopoles.

In equilibrium, a useful quasiparticle description for spin ice is in terms of deconfined magnetic charges [88]. Conversely, the long (intrinsic) lifetime of noncontractible pairs justifies their introduction as an effectively distinct “species” of quasiparticle when studying classical spin ice in the strongly nonequilibrium setting of thermal and field quenches, as demonstrated already in Refs. [302, 316].

The mere existence of noncontractible pairs in the system however does not warrant the appearance of a macroscopic metastable state. Indeed, when free monopoles are abundant, non-activated (fast) monopole-assisted decay is the leading relaxation channel with respect to thermally-activated (slow) decay of noncontractible pairs, and one does not expect any metastable plateau to appear. It is only when the system undergoes a “population inversion” (in contrast to thermodynamic equilib-

*Noncontractible pairs cannot move from the site upon which they form without separating to third neighbour distance. Otherwise, movement of the pair would require motion of a monopole along a blocked direction. Indeed, the caterpillar-like motion of separating and rejoining, trailing one another, is prevented by the intervening spin being a blocked direction for the trailing monopole.

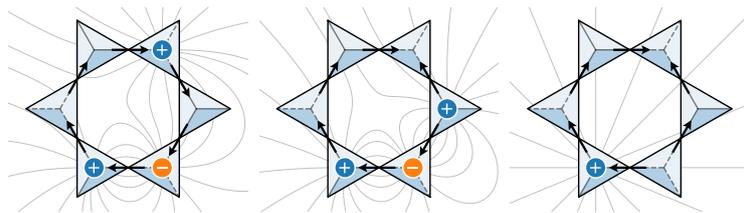


Figure 4.2: Schematic depiction of monopole-assisted decay of a noncontractible pair. A free monopole annihilates with the oppositely charged member of the stationary noncontractible pair, thereby freeing up its partner. All moves shown lower the energy of the system and hence monopole-assisted decay is the dominant decay avenue for noncontractible pairs with respect to thermally activated decay when free monopoles are abundant in the system.

rium), where noncontractible pairs become the dominant species with respect to free monopoles, that the activation energy barrier to decay can induce a long-lived metastable plateau at low temperatures. This is indeed what one observes in numerical simulations of dipolar spin ice, following appropriate thermal quenches [302].

The aforementioned population inversion is key to the metastable plateau. Its origin however was not investigated in Ref. [302] and is the subject of the present work. We find that it ultimately rests on the long-range tail of the Coulomb interaction between monopoles. This can be qualitatively understood as being due to the energetic bias in the motion of monopoles in the far field. Monopole-antimonopole collision events are subject to a Coulombic charge-charge attraction ($\propto r^{-2}$), whereas collisions between a free monopole and a noncontractible pair are subject to weaker charge-dipole interactions ($\propto r^{-3}$). This leads to a bias that increases the likelihood of free monopoles annihilating (or forming new noncontractible pairs) over their chance of annihilating existing noncontractible pairs via monopole-assisted decay. Further, since the final temperature in the thermal quenches is much less than all other energy scales in the problem, the system enters a non-hydrodynamic regime where the monopoles move at terminal velocity in the direction of the local force acting on them. This allows the system to violate the law of formal kinetics [317] and to exhibit a decay of the free monopole density faster than inverse time. The combination of the long-range bias and ‘terminal velocity’ motion of free charges leads to a rapid decay of the free monopole density in the system, leaving behind an excess of noncontractible pairs. This is ultimately the linchpin of the finite-density metastable plateau observed in numerical simulations.

This behaviour is most remarkable. By altering the dynamics of what is ultimately a transient regime, spin ice is able to enter a metastable state whose lifetime for experimentally relevant temperatures and system sizes may well exceed any realistically accessible timescales (of order one year in Fig. 4.3 when expressed in physical units).

We verify this scenario through extensive numerical Monte Carlo simulations of thermal quenches in spin ice systems with nearest-neighbour spin-spin interactions and long-range Ewald-

summed magnetic Coulomb interactions between defective tetrahedra [92]. Upon truncating the Coulomb interactions to finite range, the long-range bias is removed. We find that the finite-density metastable plateau correspondingly disappears in the thermodynamic limit. These findings are corroborated (in Sec. 4.5.2) by directly simulating mutually interacting magnetic charges hopping on a diamond lattice (with no Dirac strings), for which we observe qualitatively similar behaviour.

To supplement the numerics, we provide an analytical understanding of both the value of the plateau in the thermodynamic limit and its finite size scaling using mean field population dynamics, treating the system as a Coulomb liquid of magnetic charges. We show how the ratio of the rate of monopole-assisted decay to the rate of charge–charge annihilation underpins both the finite size scaling exponent in the case of truncated interactions, and the density at which the plateau occurs in the long-range case.

4.4 Models

In this work, we contrast the effect of truncating the Coulomb interactions between monopoles in spin ice with the same truncation in a system of magnetic charges hopping on a diamond lattice. The latter model is defined without reference to any underlying spin configuration, i.e., without Dirac strings connecting opposite charges, which allows us further clarity in ascertaining their role in thermal quenches and the formation of the metastable plateau.

4.4.1 Classical spin ice

As discussed in Sec. 1.4.3, at sufficiently low temperatures, classical spin ice (CSI) is well described by effective Ising spins $\mathbf{S}_i = S_i \mathbf{e}_i$, with $S_i \in \{-1, +1\}$, which live on the sites of the pyrochlore lattice. These spin degrees of freedom interact through the dipolar spin ice Hamiltonian

$$H_d(\{S_i\}) = \frac{J}{3} \sum_{\langle ij \rangle} S_i S_j + D \sum_{\langle ij \rangle} \left[\frac{\mathbf{e}_i \cdot \mathbf{e}_j}{|\mathbf{r}_{ij}|^3} - \frac{3(\mathbf{e}_i \cdot \mathbf{r}_{ij})(\mathbf{e}_j \cdot \mathbf{r}_{ij})}{|\mathbf{r}_{ij}|^5} \right] S_i S_j,$$

where, in the first line, we used the fact that $\mathbf{e}_i \cdot \mathbf{e}_j = -1/3$ for any nearest neighbour pair of sites, $\langle ij \rangle$.

For the majority of this work, we use the effective Hamiltonian (1.4.18) in which the exchange and dipolar interactions between the spins are retained only at nearest-neighbour level, and farther range couplings are accounted for effectively by a pairwise Coulomb interaction $V(\{Q_a\})$ between tetrahedral charges Q_a ,

$$H_c(\{S_i\}) = J_{\text{eff}} \sum_{\langle ij \rangle} S_i S_j + E_{\text{mn}} \sum_{a < b} \frac{Q_a Q_b}{r_{ab}}, \quad (4.4.1)$$

where i, j index the sites of the pyrochlore lattice, a, b index the tetrahedra and $r_{ab} = |\mathbf{r}_a - \mathbf{r}_b|/r_{\text{nn}}$ is the distance between the centres of tetrahedra a and b in units of the distance between neighbouring tetrahedra. The charge on tetrahedron a is $Q_a = \pm \sum_{i \in a} S_i/2$, where the sign depends on the sublattice that a belongs to. The charges Q_a therefore assume the values $Q_a \in \{0, \pm 1, \pm 2\}$, where $Q_a = \pm 1$ are dubbed monopoles and $Q_a = \pm 2$ double monopoles. We use the convention that a positive charge corresponds to a majority of spins pointing out of a given tetrahedron. Two equally charged monopoles on neighbouring sites have a Coulomb energy E_{nn} (in an infinite system). Throughout the chapter we use an effective exchange coupling $J_{\text{eff}} = 1.463 \text{ K}^*$ and nearest-neighbour Coulomb energy $E_{\text{nn}} = \sqrt{128/27}D = 3.06 \text{ K}$, appropriate for the classical spin ice compound $\text{Dy}_2\text{Ti}_2\text{O}_7$. Such an effective description (4.4.1) is quantitatively accurate, up to quadrupolar corrections, by virtue of projective equivalence [87] (and this is indeed the case also in thermal quenches, as illustrated in Fig. 4.5). With these parameters, the macroscopically degenerate ground state manifold corresponds to the charge vacuum, $Q_a = 0, \forall a$, i.e., a 2 in-2 out configuration of spins on each tetrahedron.

We note that the nearest-neighbour exchange interaction between spins can be viewed as a chemical potential of size $2J_{\text{eff}}$ for the monopoles (namely, the charges $Q_a = \pm 1$):

$$J_{\text{eff}} \sum_{\langle ij \rangle} S_i S_j = 2J_{\text{eff}} \sum_a Q_a^2 - N_s J_{\text{eff}}. \quad (4.4.2)$$

This interpretation however no longer holds straightforwardly in the presence of double monopoles.

To test the role of the long-range tail of the Coulomb interaction in the appearance of the population inversion, we also consider a similar model where the interactions $V(\{Q_a\})$ between monopoles are truncated at nearest-neighbour distance:

$$H_t(\{S_i\}) = J_{\text{eff}} \sum_{\langle ij \rangle} S_i S_j + \Delta \sum_{\langle ab \rangle} Q_a Q_b. \quad (4.4.3)$$

This model will be referred to as classical spin ice with truncated interactions. Such a nearest-neighbour interaction between monopoles allows for the formation of noncontractible pairs without inducing any long-range energetic bias in the motion of the monopoles.

Separating an isolated pair of nearest-neighbour monopoles with charge $Q = \pm 1$ in this model costs an energy Δ . To preserve the behaviour of the system (primarily its ground state), the truncation of the interactions must be done with care. We choose the value of Δ such that the energy barrier to separating a noncontractible pair around a hexagonal plaquette (as depicted in Fig. 4.1) is

*This value of the effective exchange coupling was obtained using the chemical potential $\mu = -8.92 \text{ K}$ in Ref. [81]. In particular, we use $\mu = -4J_{\text{eff}} - E_{\text{nn}}$ to define J_{eff} , i.e., (minus) the energy required to create a pair of oppositely charged monopoles and separate them to infinity. The dynamics of the system is however not particularly sensitive to the precise value of J_{eff} , as long as the ground state remains unchanged.

equal in the cases of truncated (4.4.3) and long-ranged (4.4.1) interactions*:

$$\Delta = E_{\text{nn}} \left(1 - \sqrt{\frac{3}{11}} \right) \simeq 1.46 \text{ K}. \quad (4.4.4)$$

Such a choice preserves the charge vacuum ground state, and since the energy barrier for the activated decay of noncontractible pairs is equal for both types of interaction, the demise of a possible metastable plateau will occur at similar times in the two cases.

The difference between the single spin flip dynamics of the two Hamiltonians, H_c and H_t , therefore rests solely in the long-range energetic bias in the motion of monopoles across the system. In a finite system containing L^3 cubic unit cells, the total number of spins is $N_s = 16L^3$, and the number of tetrahedra is $N_t = 8L^3$. In our simulations, we use periodic boundary conditions and we deal with long-ranged interactions (dipolar as well as Coulomb) using the method of Ewald summation [318, 319].

We note that there also exists a long-range Coulomb interaction between monopoles of entropic origin [320]. As we are unable to alter the range of the entropic interactions, we introduce in the following section a family of charge models that live on the diamond lattice in which the charges are not born out of underlying spin configurations. This will allow us to observe that the role of entropic interactions in thermal quenches is in fact negligible and hence they will not be discussed further in our work. This is shown most directly by the good quantitative agreement between the classical spin ice and charge model simulations, and the mean field analytics, for truncated interactions.

4.4.2 Charges on diamond lattice

To identify the role of the spin configuration underlying each monopole configuration, we also consider two further effective models of charges Q_a hopping on a diamond lattice, thereby removing any entropic effects and blocked directions associated with the spins (in particular, by removing the underlying spin network, there are no Dirac strings associated with the magnetic charges in the following models). We restrict our simulations to the relevant charge values $Q_a \in \{0, \pm 1, \pm 2\}$ only. These charge models (CM) also allow for a more direct comparison with our analytical mean field modelling (see Sec. 4.6), which largely neglects the aforementioned complications associated with the spinful description of the system dynamics.

In the case of long-range interactions between the charges, we use the Hamiltonian

$$H_c^{\text{CM}}(\{Q_a\}) = 2J_{\text{eff}} \sum_a Q_a^2 + E_{\text{nn}} \sum_{a<b} \frac{Q_a Q_b}{r_{ab}}, \quad (4.4.5)$$

*This energy barrier is equal to the Coulomb energy required to separate the pair to third neighbour distance, $r_{3\text{n}}$, i.e., $\Delta = E_{\text{nn}}(1 - r_{\text{nn}}/r_{3\text{n}})$.

subject to the hard constraint that each site may not be occupied by more than two charges. This model will be referred to as the long-range interacting charge model.

The Hamiltonian (4.4.5) must be further supplemented by rules which govern the dynamics of the charges. Namely, in order to take into account the effect of noncontractible pairs, when two opposite (single) charges come into nearest-neighbour contact, there exists some finite probability, p_{nc} , of forming a noncontractible pair. If a noncontractible pair is formed, it is then not possible for the charges to annihilate along their common bond. At finite temperature, their activated decay can be accounted for by associating an energy barrier Δ with this process.

The probability p_{nc} can be estimated by counting the number of spin configurations compatible with two oppositely charged monopoles on adjacent tetrahedra, and taking the fraction thereof that corresponds to a noncontractible pair. Considering the minimal cluster of two tetrahedra only (7 spins in total), one finds that the relevant fraction is $p_{\text{nc}} = 1/10$ [302]. Extending the calculation to larger clusters does not lead to significant variation in this value; for example, considering a full hexagon of tetrahedra involving the two monopoles gives $p_{\text{nc}} = 41/406$. Further, small perturbations in p_{nc} do not appreciably modify the dynamics of the system.

For the case of truncated interactions between the tetrahedral charges, the Hamiltonian becomes

$$H_t^{\text{CM}}(\{Q_a\}) = 2J_{\text{eff}} \sum_a Q_a^2 + \Delta \sum_{\langle ab \rangle} Q_a Q_b, \quad (4.4.6)$$

referred to as the charge model with truncated interactions. The model is again subjected to the same constraints on charge values and dynamics. The difference between the two charge models, H_c^{CM} and H_t^{CM} , lies only in the long-range energetic bias associated with the Coulomb interaction.

4.5 Monte Carlo Simulations

4.5.1 Classical spin ice

Long-range Coulomb interactions

In Fig. 4.3 we show the monopole density evolution following a thermal quench, as in Ref. [302], simulated using the modified Monte Carlo code, corresponding to (4.4.1), instead of the conventional dipolar Monte Carlo (for a direct comparison, see Fig. 4.5). We use single spin flip dynamics and the waiting time method [313, 314] introduced in Sec. 4.2.2 to access long simulation times at low temperatures. The system is initially prepared in the paramagnetic phase at infinite temperature, then at $t = 0$ the temperature is set to its target value, $T \ll J_{\text{eff}}$, and we start measuring various

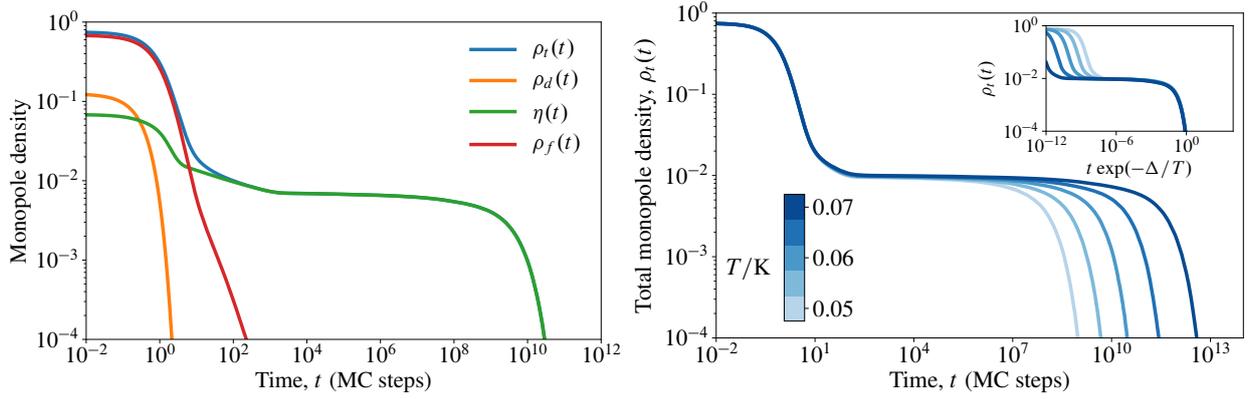


Figure 4.3: **Left:** Monte Carlo simulations of a thermal quench in spin ice subject to Ewald-summed Coulomb interactions between monopoles [Hamiltonian (4.4.1), system size $L = 22$, i.e., 170 368 spins] from infinite temperature down to $T = 0.06$ K. The curves show the evolution of the averaged total density of monopoles ρ_t (blue), the free monopole density ρ_f (red), the density of monopoles forming noncontractible pairs η (green) and the double charge density ρ_d (orange). **Right:** Total density of monopoles, ρ_t , after analogous quenches to final temperatures $T = 0.05$ – 0.07 K for a system of size $L = 8$, i.e., 8192 spins. Inset: the same curves plotted after rescaling the time axis by a factor $\exp(\Delta/T)$, where $\Delta \simeq 1.46$ K is the Coulomb energy barrier incurred by separating two monopoles around a hexagonal plaquette. In both panels, the densities are averaged over 4096 histories.

monopole densities as a function of time*. These densities are then averaged over many histories with different random initial conditions sampled from the infinite temperature ensemble. We find good agreement with the dynamical arrest observed in Ref. [302]: Rather than rapidly equilibrating to a monopole-sparse state, we observe instead the emergence of a metastable plateau in the monopole density due to noncontractible monopole–antimonopole pairs.

Specifically, we measure the total monopole density (monopoles per tetrahedron) in the system, ρ_t , counting all-in and all-out tetrahedra as doubly occupied sites; the fraction of such doubly occupied sites, ρ_d ; the density of monopoles forming noncontractible pairs, η ; and the ‘free’ monopole density[†] $\rho_f \equiv \rho_t - \eta$, i.e., the density of monopoles that do not form noncontractible pairs. A noncontractible pair is *defined* as a pair of adjacent, oppositely-charged monopoles for which the reversal of the intervening spin shared by the two tetrahedra does not lead to annihilation of the pair.

In isolation, the barrier to activated decay of a noncontractible pair is $\Delta \simeq 1.46$ K. In the presence of a finite density η of other noncontractible pairs, the distribution of energy barriers is broadened around a mean value of Δ due to dipole–dipole interactions between the pairs. Given

*In $\text{Dy}_2\text{Ti}_2\text{O}_7$ these initial conditions are experimentally relevant to initial temperatures $T \gg 1$ K, and the quench in temperature should occur over timescales less than ~ 1 ms, the characteristic single spin flip timescale [312].

[†]Note that there are many possible definitions of the ‘free’ monopole density due to ambiguities that arise in defining pairs of monopoles in the monopole-dense (short-time) limit. However, all definitions agree once the typical separation of monopoles is greater than r_{m} .

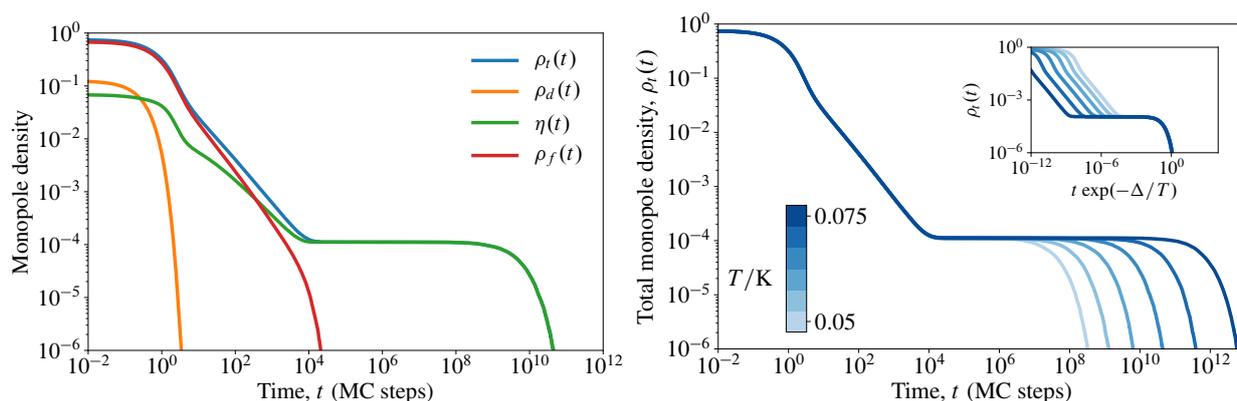


Figure 4.4: **Left:** Monte Carlo simulations of a thermal quench in spin ice where the interactions between monopoles are truncated to nearest-neighbour distance [Hamiltonian (4.4.3), system size $L = 16$, i.e., 65 536 spins] from infinite temperature down to $T = 0.06$ K. The metastable plateau due to noncontractible pairs of monopoles remains present, but occurs at lower densities and at later times than in the case of long-range interactions (cf. Fig. 4.3). **Right:** The total density of monopoles, ρ_t , after analogous quenches from infinite temperature down to final temperatures $T = 0.05$ – 0.075 K for a system of size $L = 16$, i.e., 65 536 spins. Inset: the same curves plotted after rescaling the time axis by a factor $\exp(\Delta/T)$, showing an excellent collapse of the long-time decay. In both panels time is expressed in units of Monte Carlo steps per site, and the densities are averaged over 4096 histories.

that the Coulombic approximation to the monopole–monopole interaction neglects quadrupolar corrections, we expect the distribution of such energy barriers to be more sharply peaked than in the dipolar case. This is indeed confirmed by the excellent collapse of the long-time decay of the total monopole density for various temperatures upon rescaling the time axis by a factor $\exp(\Delta/T)$, as illustrated in the right panel of Fig. 4.3 (see also Fig. 4.5, where the dipolar case shows a correspondingly broader decay of the metastable plateau).

Truncated interactions

In Fig. 4.4 we plot the various monopole densities for an identical thermal quench for the case of truncated interactions between monopoles in classical spin ice [i.e., Eq. (4.4.3)]. A metastable plateau remains present in the dynamics of the system, and once again the behaviour of the monopole densities tells us that it is clearly due to noncontractible pairs. The plateau however occurs at substantially lower densities and the onset occurs at later times when compared with the corresponding long-range interacting system, Eq. (4.4.1), in Fig. 4.3.

The decay of the monopole density at long times collapses for a range of temperatures upon rescaling the time axis by a Boltzmann factor $\exp(\Delta/T)$, as illustrated in the right panel of Fig. 4.4, confirming that the thermally activated decay of noncontractible pairs is again responsible for the eventual demise of the plateau at a time $\tau_{nc} \sim \exp(\Delta/T)$. Once a given pair has separated, the two constituent monopoles may find each other and annihilate by performing a random walk, the

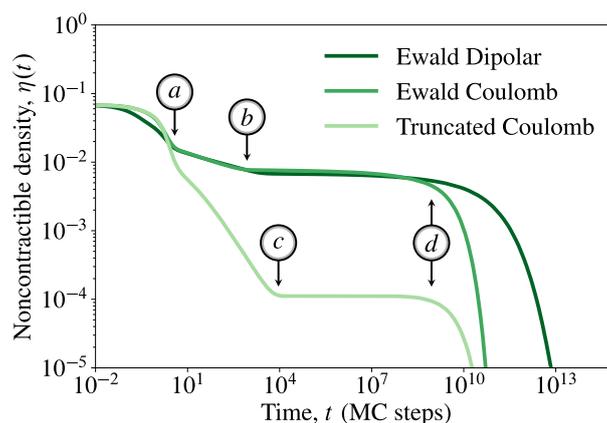


Figure 4.5: Comparison of noncontractible pair densities $\eta(t)$ for the three types of interaction for a thermal quench from infinite temperature down to $T = 0.06$ K (system size $L = 16$, i.e., 65 536 spins) in classical spin ice. Time is expressed in units of Monte Carlo steps per site, and the densities are averaged over 4096 histories. The markers labelled a , b , c , and d identify the boundaries between the four dynamical regimes discussed in the main text. At (a) , nearly all doubly occupied sites have been removed from the system. Points (b) and (c) mark the onset of the metastable plateau for the cases of long-range and truncated interactions, respectively. At (d) , the noncontractible pairs decay via thermal activation.

shortest of which is around a single hexagonal plaquette. Since the noncontractible pairs do not interact beyond a fixed, finite separation, the energy barriers are δ -distributed about Δ .

Comparison and finite size scaling

In Fig. 4.5 we plot the noncontractible pair density as a function of time, $\eta(t)$, for all three types of interaction introduced in Sec. 4.4.1 for classical spin ice: Ewald-summed dipolar interactions between spins, Ewald-summed Coulomb interactions between monopoles, and truncated (nearest-neighbour) interactions between monopoles.

In each of the three cases, the time evolution of $\eta(t)$ can be decomposed into four dynamical regimes. The crossovers between these regimes are identified by the markers (a) – (d) in Fig. 4.5.

1. At times before (a) , the large exchange energy cost associated with double occupancy of a tetrahedron ($8J_{\text{eff}}$) and the ability of such configurations to decay spontaneously ensures that doubly occupied sites are removed rapidly during this regime (exponentially fast in time, see Appendix D for further details).
2. Between (a) – (b) and (a) – (c) the differences between the three interaction types become manifest. In all cases we observe a much slower decay of the noncontractible pair density once the double monopoles have been removed from the system. However, the rate of decay and the timescales over which this decay occurs are vastly different for the truncated [(a) – (c)] versus long-range interacting [(a) – (b)] models. In the Coulomb and dipolar cases, the

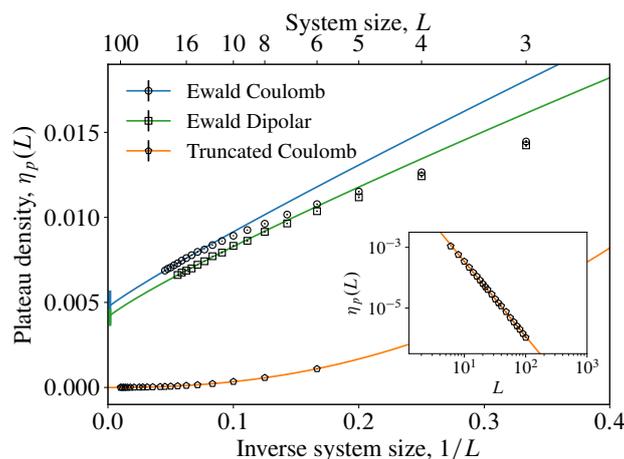


Figure 4.6: Finite size scaling of the plateau in noncontractible monopole density $\eta_p(L)$ for long-range Coulomb and truncated (nearest-neighbour) interactions between monopoles, and long-range dipolar interactions between spins in classical spin ice. The data are averaged over at least 4096 histories. The lines are fits to the scaling ansatz $\eta_p(L) - \eta_p(\infty) \sim L^{-\nu}$, while the symbols represent the Monte Carlo data. The corresponding error bars are smaller than the width of the fit lines. In the truncated case (system sizes $L = 6$ – 100 inclusive), the data are consistent with a plateau that vanishes in the thermodynamic limit. This is verified using a log–log plot of the plateau density against system size L in the inset. Conversely, the long-range Coulomb ($L = 3$ – 22 inclusive) and dipolar ($L = 3$ – 18 inclusive) cases appear to exhibit a nonvanishing noncontractible pair density in the metastable plateau in the thermodynamic limit: $\eta_p(\infty) = 4.7(9) \times 10^{-3}$ and $\eta_p(\infty) = 4.1(5) \times 10^{-3}$, respectively.

long-range nature of the interactions leads to an energetic bias which favours monopole–antimonopole (charge–charge) annihilation over monopole-assisted decay of noncontractible pairs (charge–dipole). This means that (i) the free monopoles in the system vanish more quickly, and, correspondingly, (ii) noncontractible pairs are removed more slowly than in the case of truncated interactions. Since the plateau forms when there are no free monopoles left in the system, point (i) gives rise to the earlier onset of the plateau, while point (ii) implies that the plateau forms at a higher density.

3. The metastable plateau occurs between (b)–(d) and (c)–(d). This regime, in which the system contains essentially only noncontractible pairs, spans many orders of magnitude in time at the low final quench temperatures considered in this chapter.
4. At times after (d), noncontractible pairs are able to decay via thermal activation, leading to the demise of the metastable plateau. This occurs at a time $\tau_{nc} \sim \exp(\Delta/T)$.

By construction, the decay of the plateau occurs at similar times for the models with truncated and long-range Coulomb interactions between monopoles. The difference in the decay times between the Coulomb and dipolar models is due to the larger variance in energy barriers for activated decay of the pairs in the latter. Indeed, one may model the decay of the plateau by assuming a

Gaussian distribution of energy barriers, $P(\epsilon)$, with mean Δ and width σ . The activated decay of the noncontractible pair density $\eta(t)$ is then approximated as $\eta(t) = \int d\epsilon P(\epsilon)e^{-t/\tau(\epsilon)}$, where the decay time $\tau(\epsilon) \propto e^{\epsilon/T}$. The values $\sigma_d \simeq 0.1$ K [302], $\sigma_c \simeq 0.03$ K and $\sigma_t \simeq 0$, lead to the best fit of the Monte Carlo data (not shown).

Notice that, in systems of finite size, the appearance of a noncontractible plateau in the averaged monopole density is, in fact, unavoidable. On the one hand, the probability that all free monopoles annihilate before all noncontractible pairs have decayed is finite; and, if this happens, the only decay process left for the noncontractible pairs is activated decay. On the other hand, even when the last two monopoles in the system are free, there exists a finite probability of forming a new noncontractible pair, rather than annihilation, when the two monopoles come into nearest-neighbour contact. The latter process places a hard, nonzero lower bound on the density of the noncontractible plateau of $O(1/L^3)$, which is purely a finite size effect.

In order to understand the origin of the plateau and the difference in behaviour between the truncated and long-range interacting models, we ought therefore to look at the finite size scaling behaviour of the plateau density. Figure 4.6 shows the noncontractible monopole density in the plateau, $\eta_p(L)$, for systems of different sizes (parameterised by the linear system size L) and the same final quench temperature $T = 0.06$ K. We perform a fit to the scaling ansatz $\eta_p(L) - \eta_p(\infty) \sim L^{-\nu}$, to extract the exponent ν , the value of the plateau in the thermodynamic limit, $\eta_p(\infty)$, and the constant of proportionality. The form of this scaling ansatz is justified later in Sec. 4.6, where we show that a power law decay of the free monopole density with time implies power law scaling of the metastable plateau density with system size. Hence, the scaling ansatz only applies once any transient (non-power law) behaviour of $\rho_f(t)$ at short times has subsided. For dipolar interactions between spins, it is not numerically feasible to access system sizes sufficiently large to observe an asymptotic power law decay regime of the free monopole density. We nevertheless provide a fit to the data in this case, but it should be noted that the resulting parameters are subject to some degree of systematic error. In the case of Coulomb interactions between the monopoles, such asymptotic power law decay of $\rho_f(t)$ is observed in systems of size $L \geq 14$ (i.e., 43 904 spins), and correspondingly only these data are included in the scaling analysis.

The inset of Fig. 4.6 demonstrates that the metastable plateau in the truncated case is indeed a finite size effect: The number of noncontractible pairs in the plateau increases subextensively with the size of the system, $\nu \simeq 2.46$, and the density $\eta_p(\infty)$ is consistent with a vanishing value in the thermodynamic limit. By contrast, in the case of long-range interactions, the number of noncontractible pairs in the plateau scales *extensively* with system size, with subleading, subextensive contributions. Hence, the density of the plateau in the long-range case tends asymptotically towards a finite value, also shown in Fig. 4.6. The subextensive corrections give rise to the L -dependence of the plateau density. The finite size scaling exponent in this case is $\nu = 0.9(3)$.

We shall summarise these results and attempt to understand the origin of the different behaviours

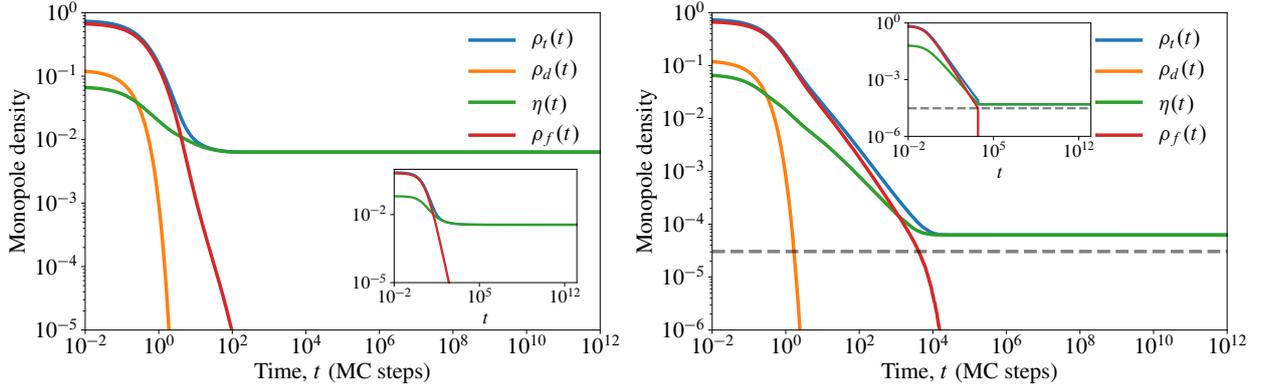


Figure 4.7: **Left:** Monte Carlo simulations of charges hopping on the diamond lattice subject to long-range Coulomb interactions [Hamiltonian (4.4.5), system size $L = 22$, i.e., 170 368 spins] from infinite temperature down to zero temperature. The analytic solution to the mean field equations for the charge densities is shown in the inset for comparison. **Right:** Monte Carlo simulations of charges hopping on the diamond lattice subject to truncated (nearest-neighbour) interactions [Hamiltonian (4.4.6), system size $L = 16$, i.e., 65 536 spins] from infinite temperature down to zero temperature. The analytic solution, (4.6.6), to the mean field equations for the charge densities is shown in the inset for comparison. The dashed lines indicate the threshold density corresponding to the disappearance of free charges in a system of finite size, $\rho_* = 1/N_t$. In both panels time is expressed in units of Monte Carlo steps per site, and the data are averaged over 4096 histories.

and exponents by modelling the time evolution of the system using mean field population dynamics in Sec. 4.6.

4.5.2 Charges on diamond lattice

Long-range Coulomb interactions

Moving to the charge description, characterised by the long-range charge model Hamiltonian (4.4.5), $H_c^{\text{CM}}(\{Q_a\})$, we obtain the results shown in Fig. 4.7 for a thermal quench down to zero temperature. As long as the final quench temperature satisfies $T \lesssim E_{\text{nn}}/L^2$, the dominant effect of changing temperature is to modify the long-time activated decay of the plateau. We therefore focus on the limit of zero temperature for simplicity.

The initial distribution of the charges is set using an infinite temperature distribution of spins on the bonds of the diamond lattice, i.e., using the same initial conditions as in Sec. 4.5.1. After initialisation of the system, all reference to an underlying spin configuration is removed, and the time evolution is determined by the dynamical rules laid out in Sec. 4.4.2. The most significant difference therefore between the charge model and spin ice systems is the blocked directions imposed by the spins in the latter. As in the case of the spinful simulations, we measure the various monopole densities as functions of time after the thermal quench and average over histories.

In this case, we observe a plateau that occurs at finite density and which persists indefinitely since the noncontractible pairs cannot undergo activated decay at zero temperature. However, contrasting Figs. 4.3 and 4.7, there are some quantitative differences between the dynamics of the charge and the spin models. In particular, the decay of free monopoles occurs much more quickly in the charge model given the same type of interactions. This implies that the onset of the plateau occurs significantly earlier in time than the corresponding model in CSI (cf. Fig. 4.3).

Truncated interactions

As shown in the right panel of Fig. 4.7, in the case of truncated interactions between charges, Eq. (4.4.6), we again observe a plateau that occurs at later times and at lower densities than in the long-range interacting charge model (Fig. 4.7). The free charge density decays approximately as $1/t$ in the long-time limit, i.e., after the double charges have been removed from the system, while the noncontractible pair density also decays as a power law in time, but with a smaller exponent. The power law decay of these quantities is cut off when the free monopoles drop below $O(1/L^3)$ density, as indicated by the dashed line in Fig. 4.7. The noncontractible pairs that remain in the system can only further decay by thermal activation and the noncontractible plateau is thus established when the free monopole density crosses this threshold.

Comparison and finite size scaling

The finite size scaling of the plateau in the case of charges hopping on the diamond lattice, contrasting the behaviour of Eqs. (4.4.5) and (4.4.6), is presented in Fig. 4.8. We again observe that the long-range interacting case tends towards a finite plateau density in the thermodynamic limit, while the plateau is merely a finite size effect in the case of truncated interactions between the charges, i.e., $\lim_{L \rightarrow \infty} \eta_p(L) = 0$ with $\nu \simeq 2.28$.

These findings corroborate the conclusions of Sec. 4.5.1 pertaining to classical spin ice. In particular, that the plateau is not a finite size effect in the case of long-range Coulomb interactions between charges. Since the subleading corrections decay more quickly in the charge description, $\nu = 1.8(4)$, we are able to make this claim on even stronger terms.

The fact that the finite size scaling of the plateau, i.e., the exponent ν , differs significantly between the spinful and charge descriptions for the long-range case, while it is very similar between the spinful and charge descriptions for truncated interactions, is a puzzle that we shall attempt to understand in Section 4.6. Indeed, we will see that one can achieve a great deal of analytical insight into the observed behaviour by means of appropriate mean field modelling.

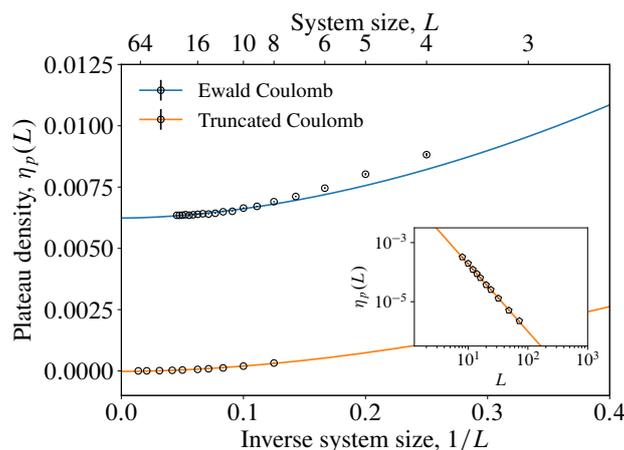


Figure 4.8: Finite size scaling of the noncontractible plateau density $\eta_p(L)$ for the case of charges hopping on the diamond lattice subject to long-range Coulomb and truncated interactions. The data are averaged over at least 4096 histories. The lines are fits to the scaling ansatz $\eta_p(L) - \eta_p(\infty) \sim L^{-\nu}$, while the symbols represent the Monte Carlo data. The corresponding error bars are smaller than the width of the fit lines. As in CSI, the case of truncated interactions ($L = 6\text{--}72$ inclusive) is consistent with a vanishing plateau density in the thermodynamic limit, verified by the log–log plot of plateau density against linear system size in the inset. Conversely, the long-range Coulomb case ($L = 4\text{--}22$ inclusive) exhibits a nonvanishing plateau density in the thermodynamic limit: $\eta_p(\infty) = 6.24(2) \times 10^{-3}$.

4.6 Summary and mean field modelling

From our simulations we see that the behaviour of the four models in question is visibly similar. The key differences are: (i) the finite size scaling of the plateau is consistent with a finite versus a vanishing value in the thermodynamic limit in the case of long-range versus truncated interactions, respectively, both in CSI and the CM; moreover, in the case of long-range interactions, (ii) the decay of $\rho_f(t)$ is notably faster, and the variation with system size L is stronger (i.e., ν is significantly larger), in the charge simulations than in the spin ice simulations.

Regarding the discrepancy in the decay of the free monopole density, highlighted in point (ii) above, the most significant difference between the dynamics of the two models in the regime where monopoles are sparse is the existence of blocked directions in classical spin ice*. That is, for a given (isolated) free monopole, there always exists one of four directions (corresponding to the minority spin) along which the monopole cannot hop, as shown schematically in Fig. 4.9. Assuming that the direction of the local Coulomb field is distributed randomly over the unit sphere, the fraction of charges which are unable to lower their energy due to blocking is $\Omega_b/4\pi$, where Ω_b is the solid angle for which there is a positive projection onto exactly one of the local basis vectors $\{\mathbf{e}_i\}$. This

*One may also wonder whether the differences in the short-time dynamics affect significantly the asymptotic decay of the free monopole density. This has been ruled out by changing between the dynamics generated by long-range classical spin ice and the long-range charge model at some later time, say $t = 10$ MC steps (data not shown).

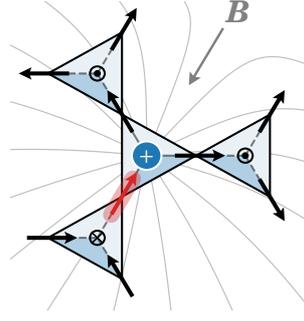


Figure 4.9: Schematic depiction of a blocked direction for a free monopole. The isolated monopole is unable to move along the bond containing the minority spin (shown in red), since its reversal would lead to the creation of a double charge on the central tetrahedron. Blocked directions have a significant quantitative impact on the dynamics of monopoles subject to long-range interactions by instantaneously pinning some finite fraction of free monopoles. The local magnetic field \mathbf{B} determines which direction(s) lower the energy of the system; if this direction is unique and coincides with the blocked direction (as in the figure), then the monopole is pinned and cannot move along any direction without overcoming a finite energy barrier.

leads to a probability

$$p_b = \frac{\Omega_b}{4\pi} = \frac{3}{2\pi} \left[\frac{\pi}{3} - \arctan \sqrt{2} \right] \simeq 4.4\%, \quad (4.6.1)$$

for a given free monopole to be pinned (at zero temperature) due to blocking. In addition, even when the monopole is not pinned, the available phase space for motion is reduced by blocking. Notice that (4.6.1) underestimates the effect of pinning, because at the lattice scale the direction of the Coulomb interaction is correlated with the bond directions, which violates the assumption of uniformity over the unit sphere. Hence, we conclude that a finite fraction of monopoles, lower-bounded by (4.6.1), are instantaneously* pinned in the spinful description due to the interplay of interactions and blocked directions. It is then reasonable to expect that the free monopole density decays more slowly in the presence of such pinned charges. While this is an interesting aspect of stochastic processes in spin ice that warrants further investigation (maybe by including some effective disorder in the relevant equations governing the dynamics of the charges), it is beyond the scope of the present work. We shall nonetheless see below that this effect plays a key quantitative role in the difference between long-range CSI and CM results.

The scaling fits to the Monte Carlo data $\eta_p(L) - \eta_p(\infty) \sim L^{-\nu}$ give the values summarised in Tab. 4.1. In the following, we show how one can understand this behaviour qualitatively and

*Since the spatial distribution of monopoles changes with each Monte Carlo step, the effect of pinning is transient—those monopoles which are pinned at one time may later become unpinned depending on the distribution of free monopoles. Estimating the relevant timescale is generally difficult, but one may expect that changing the angle of the local force acting on a given monopole requires a rearrangement of the spatial distribution of free monopoles on the order of their typical separation, which takes a characteristic time $\sim \rho^{-1/3}$.

Model	Interactions	Plateau value, $\eta_p(\infty)$	Scaling exponent, ν
CSI	truncated	0	2.46(1)
	long-range	$4.7(9) \times 10^{-3}$	0.9(3)
CM	truncated	0	2.28(2)
	long-range	$6.24(2) \times 10^{-3}$	1.8(4)

Table 4.1: Summary of finite size scaling results for both systems and both types of interaction between the tetrahedral charges. The scaling ansatz $\eta_p(L) - \eta_p(\infty) \sim L^{-\nu}$ was used to obtain the values shown in the table.

sometimes even quantitatively using mean field population dynamics of reaction diffusion processes. This allows us to model the time evolution of the monopole/charge densities and to obtain estimates of the finite size scaling exponents to compare with our numerical results.

4.6.1 Short-time dynamics

If we want to describe the simulations in terms of reaction-diffusion processes between (effective) particles, we ought to consider in principle five different species in addition to the charge vacuum: positive and negative single and double charges, and noncontractible pairs. The noncontractible pairs are immobile, pinned to the bond on which they form, and can undergo the activated and monopole-assisted decay processes discussed previously. Single charges are able to move freely throughout the lattice (neglecting the effects of spin blocking/pinning). The double charges can either decay spontaneously into two single charges of the same sign, if adjacent to an empty site, or they can be hit by a single charge of the opposite sign and decrease their charge by one, thus producing a single (mobile) charge. Finally, two adjacent double charges of opposite sign can decay to form a noncontractible pair by flipping the intervening spin. All decay processes involving double charges reduce the energy of the system, and thence are able to occur spontaneously, even at zero temperature.

The rate of decay of double monopoles does depend on the free monopole density; however it is easy to convince oneself that the ‘phase space’ for decay (either spontaneous or monopole-assisted) is always larger than that for processes which preserve the number of double charges, and it becomes progressively more so as the free monopoles decay in time. Their evolution thus rapidly decouples from the other species and becomes exponentially fast in time: asymptotically $\rho_d(t) \propto e^{-7t/2}$, as argued in Appendix D, which appears to fit well all simulations.

The single charges that are produced in the decay of double charges merely become a known time-dependent source term in the corresponding equation governing their density; as we see from the simulations, this contribution becomes irrelevantly small for $t \gtrsim 1$. When looking at the total or

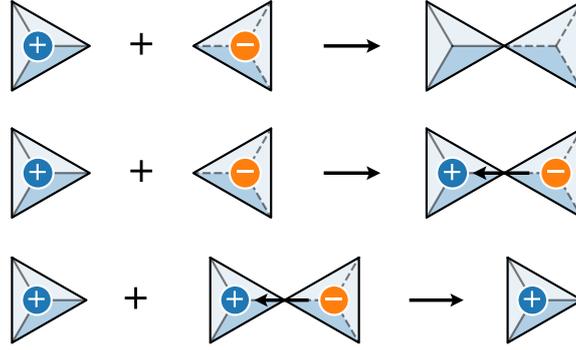


Figure 4.10: Schematic illustration of the processes that enter the reaction–diffusion equations. From top to bottom: (i) oppositely charged free monopoles may annihilate to give the vacuum, (ii) two oppositely charged free monopoles may combine to form a noncontractible pair, and (iii) a free monopole and a noncontractible pair can undergo monopole-assisted decay.

free monopole/charge densities, the double charges contribute towards the ‘hump’ observed at short times, before the onset of the asymptotic power-law behaviour. It can be shown that the double charge contribution indeed does not affect the asymptotic scaling behaviour we are interested in understanding, affecting only the density of the noncontractible plateau.

For these reasons, in the following, we shall ignore the double charges altogether and focus on the three remaining species of particle: positively and negatively charged free monopoles living on the sites of a diamond lattice, with densities $\rho_q(t)$ (charge $q = \pm$); and immobile noncontractible pairs living on the bonds, with density $\eta(t)$. The equations determining their dynamics are presented and analysed in the following sections.

4.6.2 Truncated interactions

The mean field equations (i.e., neglecting spatial fluctuations) describing the time evolution of the monopole densities in the case of truncated interactions between monopoles are (for $q = \pm$)

$$\frac{d\rho_q}{dt} = -\mathcal{K}\rho_+\rho_-, \quad (4.6.2)$$

$$\frac{d\eta}{dt} = -\frac{\mathcal{R}}{2}(\rho_+ + \rho_-)\eta + \mathcal{K}'\rho_+\rho_-. \quad (4.6.3)$$

The three processes that contribute towards these coupled equations are shown schematically in Fig. 4.10. Equation (4.6.2) describes the annihilation of oppositely charged free monopoles, which occurs with rate \mathcal{K} . The first term in Eq. (4.6.3), with rate \mathcal{R} , describes the monopole-assisted decay of noncontractible pairs—a free monopole annihilates the member of a noncontractible pair with the opposite sign. Such a process removes two monopoles previously forming a noncontractible pair, but preserves the number of free monopoles in the system, and therefore does not appear

in (4.6.2). Finally, the second term in (4.6.3) describes the probabilistic formation of noncontractible pairs when two oppositely charged monopoles come into nearest-neighbour contact. As we want to understand the origin and scaling behaviour of the noncontractible pair plateau, we are not interested in the very long-time behaviour of the system. We have therefore disregarded the terms corresponding to the activated decay of the noncontractible pairs. Equivalently, (4.6.2) and (4.6.3) describe the zero-temperature dynamics of the system.

Charge neutrality ensures that $\rho_+(t) = \rho_-(t)$ for all times, allowing us to solve (4.6.2) for the time evolution of the free monopole densities $\rho_q(t)$:

$$\rho_q(t) = \frac{\rho_q^0}{1 + \mathcal{K}\rho_q^0 t}, \quad (4.6.4)$$

where $\rho_q^0 \equiv \rho_q(0)$. This solution may then be substituted into (4.6.3) describing the noncontractible monopole density $\eta(t)$

$$\frac{d\eta}{dt} + \mathcal{R}\rho_q(t)\eta = \mathcal{K}'\rho_q^2(t), \quad (4.6.5)$$

which can also be solved exactly to give

$$\eta(t) = \frac{(\mathcal{K}'/\mathcal{K})\rho_q^0}{(\mathcal{R}/\mathcal{K} - 1)(1 + \mathcal{K}\rho_q^0 t)} + \left[\eta_0 - \frac{(\mathcal{K}'/\mathcal{K})\rho_q^0}{\mathcal{R}/\mathcal{K} - 1} \right] \frac{1}{(1 + \mathcal{K}\rho_q^0 t)^{\mathcal{R}/\mathcal{K}}}. \quad (4.6.6)$$

Evidently, the long-time behaviour of the noncontractible monopole density $\eta(t)$ depends crucially on the ratio of rate constants \mathcal{R}/\mathcal{K} . If $\mathcal{R}/\mathcal{K} < 1$, then the second term in (4.6.6) dominates at long times and the noncontractible pairs decay more slowly than the free monopoles, as is observed in the numerics, illustrated in particular in Fig. 4.4 (this is also consistent with the analytic estimates of \mathcal{R}/\mathcal{K} that we present below).

In the thermodynamic limit, these equations predict that there is no plateau in the noncontractible pair density since both $\rho_q(t)$ and $\eta(t)$ may become arbitrarily small. However, in a system of finite size containing $8L^3$ tetrahedra, the decay of $\rho_q(t)$ is cut off when the free monopole density reaches $O(1/L^3)$: $\rho_q(t_*) \sim L^{-3}$, i.e., at a time $t_* \sim L^3$ corresponding to the removal of *all* free monopoles in a finite system. If the noncontractible pair density decays more slowly, as is the case for $\mathcal{R}/\mathcal{K} < 1$, there is still a finite density of noncontractible pairs present in the system at t_* , and they can further decay only via thermal activation. The value of this density scales as $\eta(t_*) \sim t_*^{-\mathcal{R}/\mathcal{K}}$ for sufficiently large $t_* \gg (\mathcal{K}\rho_q^0)^{-1}$, allowing us to deduce the leading order term in the dependence of the noncontractible plateau on system size:

$$\eta(t_*) \sim L^{-3\mathcal{R}/\mathcal{K}}, \quad (4.6.7)$$

and therefore extract the exponent $\nu = 3\mathcal{R}/\mathcal{K}$.

We can estimate the ratio \mathcal{R}/\mathcal{K} from the microscopic details of our system as the product of two contributions,

$$\frac{\mathcal{R}}{2\mathcal{K}} = \frac{N_{\mathcal{R}}}{N_{\mathcal{K}}} \cdot \frac{\tau_{\mathcal{K}}}{\tau_{\mathcal{R}}} \simeq \frac{3}{4} \cdot \frac{1}{2}. \quad (4.6.8)$$

The first factor in (4.6.8), $N_{\mathcal{R}}/N_{\mathcal{K}}$, comes from the fact that a free monopole has 4 adjacent free legs along which another free monopole may approach, while a noncontractible pair has only 3 (one of the four total legs being blocked by the other member of the pair)*. Therefore the factor $3/4$ encodes the relative sizes of the basins of attraction in the two cases. The second factor $\tau_{\mathcal{K}}/\tau_{\mathcal{R}}$ derives from the ratio of timescales—in the case where two free monopoles are approaching one another, both are mobile, while in the case of a free monopole approaching a noncontractible pair, the noncontractible pair is pinned and only the free monopole is mobile. This leads to a factor of 2 difference in the (random walk) timescales for the two processes. The factor of $1/2$ on the left hand side of (4.6.8) originates from the definition of \mathcal{R} in (4.6.3). We therefore estimate that $\mathcal{R}/\mathcal{K} \simeq 3/4$, and correspondingly the noncontractible plateau scales approximately as

$$\eta(t_*) = \eta_p(L) \sim L^{-9/4}, \quad (4.6.9)$$

in the case of truncated interactions between charges.

This estimate can be improved upon by examining larger clusters. Indeed, including next-nearest neighbours in the cluster, the presence of blocked directions leads to a small correction to the finite size scaling exponent in the case of CSI, as shown in Appendix C, while it remains unchanged for the CM:

$$\nu_{\text{CSI}} \simeq \frac{90}{37} \simeq 2.43, \quad \nu_{\text{CM}} \simeq \frac{9}{4} = 2.25. \quad (4.6.10)$$

These approximate exponents are consistent with the values $\nu = 2.46(1)$ and $\nu = 2.28(2)$ obtained from the Monte Carlo data in Figs. 4.6 and 4.8, respectively. Note that the absolute values of \mathcal{R} and \mathcal{K} differ substantially between CSI and the CM due to the presence of blocked directions in the former, but their *ratio* remains essentially the same.

We are now able to understand why the spinful and charge descriptions exhibit quantitatively similar behaviour. In both cases, the charges exhibit diffusive motion (until they become nearest neighbours, at which point they deterministically annihilate). The numerical results suggest that the annealed (random) blocked directions do not significantly affect the diffusive motion of the charges, and therefore do not alter the form of the decay of the free monopole density. This is because the motion of monopoles across the system (i.e., beyond nearest-neighbour separation) is not subject to any energetic bias controlling the direction of their motion. Hence, the insertion of blocked directions at random has little effect on the purely random motion of charges when averaged over

*We note that in the spinful description, the rate constant \mathcal{K} includes the formation of noncontractible pairs in addition to annihilation events, and so blocked directions do not alter this argument to leading order.

histories—no monopoles are instantaneously pinned due to blocking. This is also evidence of the fact that entropic interactions in CSI due to the underlying spins do not play a significant role in the evolution of the monopole density following a thermal quench. The free monopole density decays as $1/t$ in both CSI and the CM with truncated interactions, and we consequently obtain a vanishing plateau in the thermodynamic limit. Further, the value of ν is set by the ratio of the rates of monopole-assisted decay to free monopole annihilation, which is common to both descriptions, up to small corrections which result from the impact of blocked directions on the microscopic annihilation process.

4.6.3 Long-range Coulomb interactions

In Sec. 4.6.2 we were able to develop a rather complete understanding of the case of truncated interactions, which largely hinged on the $1/t$ scaling of the free monopole density. We would now like to study how the behaviour changes in the presence of long-range interactions. One could naïvely try to introduce them at the level of the reaction diffusion equations; however, this is known to recover the law of formal kinetics at long times, i.e., $1/t$ behaviour of $\rho_f(t)$, which leads to the same conclusion of a vanishing plateau value in the thermodynamic limit. This is however in contradiction with the observation that $\rho_f(t)$ decays faster than $1/t$ in our Monte Carlo simulations of long-range interacting systems (and with the observation of a finite value for the metastable plateau).

As is often the case, the devil lies in the details. In order to observe a long-lived metastable plateau, we need to quench to very low temperatures, $T \ll J_{\text{eff}}$. In a discrete system with long-range interactions and finite lattice spacing, the hydrodynamic description of Refs. [317] and [321] does not always apply to the Monte Carlo time evolution of our simulations. Take for example the limiting case of a quench to zero temperature. The quasiparticles move only downwards or across in energy, $\delta E \leq 0$, and they move at ‘terminal velocity’ (i.e., one lattice spacing per unit time) irrespective of the strength of the force acting upon them. On the contrary, the hydrodynamic description applies when the Monte Carlo process is a (lightly) biased random walk, $|\delta E| \ll T$, and the equations of motion approximately take the familiar overdamped form where the velocity of the particles is proportional to the force acting on them. This is how our simulations violate the law of formal kinetics (at intermediate times) and achieve a decay of free monopole density which is faster than $1/t$ at the low temperatures studied in this chapter.

Modelling the strictly-biased motion at terminal velocity is a tall order. However, at mean field level, one can put forward the following approximate argument: the free monopole density decays with a time constant given by the time taken to travel at terminal velocity to the next free monopole, some characteristic distance $\rho^{-1/d}$ away, namely $\tau_{\text{iv}} \sim \rho^{-1/d}$, where d is the dimensionality of the

system. Then we have

$$\frac{d\rho}{dt} \propto -\frac{\rho}{\tau_{\text{tv}}} \quad \Rightarrow \quad \rho(t) \sim 1/t^d. \quad (4.6.11)$$

This behaviour is in very good agreement with the $\rho_f(t)$ decay observed in the CM with long-range interactions if one neglects the formation of noncontractible pairs. We shall delay the discussion of the CSI case to later in this section.

In the absence of long-range interactions, there are no forces beyond a fixed finite separation between monopoles and they perform an unbiased random walk, even at zero temperature. It then takes a characteristic time, ρ^{-1} , corresponding to the time taken for a monopole to explore its characteristic volume in three dimensions, to come in contact and annihilate with another monopole. In this case, τ_{tv} should be replaced by $\tau_{\text{tw}} \sim \rho^{-1}$ and one recovers the $1/t$ scaling obtained more rigorously in Sec. 4.6.2.

In order to express all these considerations more formally, and to take into account explicitly the noncontractible pair density $\eta(t)$, which has been ignored thus far, it is convenient to introduce the following phenomenological reaction diffusion equations

$$\frac{d\rho_q}{dt} = -\mathcal{K}[\rho_+(t)\rho_-(t)]^{(1+\beta)/2}, \quad (4.6.12)$$

$$\frac{d\eta}{dt} = -\frac{\mathcal{R}}{2}(\rho_+ + \rho_-)\eta - \frac{\mathcal{K}'}{\mathcal{K}} \frac{d\rho_q}{dt}, \quad (4.6.13)$$

with the parameter $\beta \leq 1$ (with $\beta = 1$ corresponding to the truncated case, and $\beta = 1/3$ corresponding to the terminal velocity argument given above, neglecting the effect of nonzero η)*.

Using charge neutrality $\rho_+(t) = \rho_-(t)$, the first of these equations gives rise to a free monopole density

$$\rho_q(t) = \frac{\rho_q^0}{(1 + \beta\mathcal{K}_0\rho_q^0 t)^{1/\beta}}, \quad (4.6.14)$$

where we have defined for convenience $\mathcal{K}_0 \equiv \mathcal{K}(\rho_q^0)^{\beta-1}$. The parameter β sets the asymptotic rate of decay of the free monopole density in the system: $\rho_q(t) \sim t^{-1/\beta}$. This decay is faster than the truncated case ($\rho_q \sim 1/t$) when $\beta < 1$. Defining

$$\Theta(t) = \int_0^t dt' \rho_q(t') \quad (4.6.15)$$

$$= \frac{1}{1-\beta} \frac{1}{\mathcal{K}_0} \left[1 - (1 + \beta\mathcal{K}_0\rho_q^0 t)^{(\beta-1)/\beta} \right], \quad (4.6.16)$$

*We have defined β in this way in order to make the formulae that follow neater and more compact.

the solution for the noncontractible monopole density may be written as

$$\eta(t) = e^{-\mathcal{R}\Theta(t)} \left[\eta_0 + \int_0^t dt' e^{\mathcal{R}\Theta(t')} \mathcal{K}' [\rho_+(t') \rho_-(t')]^{(1+\beta)/2} \right]. \quad (4.6.17)$$

It is possible to obtain an analytic expression for $\eta(t)$ by expressing the integral in (4.6.17) in terms of the incomplete Gamma function. Since, for $\beta < 1$, $\Theta(t)$ tends towards a constant at large times, the solution for $\eta(t)$ exhibits a plateau at finite density, $\eta(t) \rightarrow \eta_\infty$, as $t \rightarrow \infty$. The density at which this plateau occurs is

$$\eta_\infty = e^{-\alpha\mathcal{R}/\mathcal{K}_0} \left\{ \eta_0 + \rho_q^0 \frac{\alpha\mathcal{K}'}{\mathcal{K}} e^{\alpha\mathcal{R}/\mathcal{K}_0} \left[\frac{\alpha\mathcal{R}}{\mathcal{K}_0} \right]^{-\alpha} \gamma \left(\alpha, \frac{\alpha\mathcal{R}}{\mathcal{K}_0} \right) \right\}, \quad (4.6.18)$$

where $\alpha \equiv 1/(1-\beta)$, and $\gamma(s, x)$ is the lower incomplete gamma function. Hence, the value of the plateau is exponentially sensitive to the ratio of rate constants $\mathcal{R}/\mathcal{K}_0$, and vanishes as $\beta \rightarrow 1^-$ (i.e., $\alpha \rightarrow \infty$).

At sufficiently large times,

$$\eta(t) \simeq \eta_\infty \left[1 + \frac{\alpha\mathcal{R}}{\mathcal{K}_0} (\beta\mathcal{K}_0\rho_q^0 t)^{(\beta-1)/\beta} \right]. \quad (4.6.19)$$

The finite size scaling of the noncontractible plateau then follows from the fact that the free monopole decay is cut off at a time t_* , defined by $\rho_q(t_*) \sim L^{-3}$. As before, t_* equals the time at which free monopoles are completely removed from a system of finite size. This gives $t_* \sim L^{3\beta}$ and correspondingly the finite size scaling of the plateau satisfies

$$\eta(t_*) - \eta_\infty \sim t_*^{-(1-\beta)/\beta} \sim L^{-3(1-\beta)}. \quad (4.6.20)$$

The scaling exponent of the plateau, ν , can therefore be directly related to the exponent β which quantifies the asymptotic rate of decay of the free monopole density,

$$\nu = 3(1-\beta). \quad (4.6.21)$$

This relationship is consistent with the discrepancy between the finite size scaling exponents in the long-range interacting CSI and CM cases: The rapid decay of the free monopole density permitted by the lack of blocked directions in the CM case implies a larger β^{-1} and, hence, a larger ν . Indeed, numerically fitting the exponent of the asymptotic free monopole decay, we obtain $\beta^{-1} \simeq 1.4$ and $\beta^{-1} \simeq 2.3$ corresponding, through (4.6.21), to scaling exponents $\nu \simeq 0.86$ and $\nu \simeq 1.7$ for the cases of long-range CSI and the CM, respectively. These values are in reasonable agreement with those obtained from the numerical finite size scaling analysis: $\nu = 0.9(3)$ and $\nu = 1.8(4)$.

Notice that the mean field equations (4.6.12) and (4.6.13) can only be expected to hold at asymptotically long times for zero-temperature quenches. For any finite T , as the monopoles become sparser, the forces between them become weaker and eventually one reaches the hydrodynamic regime, $|\delta E| \ll T$, discussed earlier, and a $1/t$ decay of $\rho_f(t)$ ensues. The typical Coulomb interaction felt by a given monopole through the separation $\rho(t)^{-1/d}$ is (in $d = 3$ for concreteness)

$$\langle E_c(t) \rangle \sim -E_{\text{nn}}\rho(t)^{1/3}. \quad (4.6.22)$$

The corresponding *change* in Coulomb energy when moving a free monopole to an adjacent site then scales as

$$\langle \delta E_c(t) \rangle \sim E_{\text{nn}}\rho(t)^{2/3}. \quad (4.6.23)$$

Assuming $\rho(t) \sim 1/t^{1/\beta}$, the time threshold $\langle \delta E_c(t) \rangle \sim T$ corresponding to the crossover to $1/t$ decay of ρ_f can then be estimated to scale with temperature as $t_T \sim (E_{\text{nn}}/T)^{3\beta/2*}$. The crossover can be observed in our Monte Carlo simulations at sufficiently high temperatures; however it is barely visible within the accessible system sizes and the corresponding plots are not very informative, and we refrain from showing them here. From (4.6.13), we deduce that the noncontractible plateau therefore begins to decay at times $t \gtrsim t_T$. The rate of decay however vanishes as temperature is lowered, i.e., $\ln \eta \sim -T^{v/2} \ln t$. The zero-temperature limit therefore does not commute with the limit of infinite time. If the latter is taken first, the plateau decays to a vanishing thermodynamic value at large times. If the former is taken first, then a finite plateau survives. Since the timescale for activated decay of the plateau scales exponentially with temperature, while t_T scales algebraically [at least for a power law decay of $\rho_f(t)$], it will be the case that $t_T < \exp(\Delta/T)$ at the low but nonzero quench temperatures that we considered in this chapter. For systems of finite size, the relevant question then becomes whether t_T is larger or smaller than the time t_* that it takes for the free monopole density to become less than $O(1/L^3)$.

We finally note that even at zero temperature the mean field equations will eventually break down at a time corresponding to single charge densities ρ_q at which free charges become so dilute that the bias for free charge–charge annihilation over monopole-assisted decay is removed. We term such a time t_d , which may be obtained by comparing $\langle \delta E_c(t) \rangle$ with the typical energy due to charge–dipole interactions with the noncontractible pairs present in the metastable plateau regime. Once this bias disappears, monopole-assisted decay may once again become favourable and the plateau is able to gradually decay.

The phenomenological model that we have presented illustrates in a simple manner the mecha-

*Taking the terminal velocity limit, $\beta = 1/3$, and using the parameters for $\text{Dy}_2\text{Ti}_2\text{O}_7$, the thermal crossover occurs at $t_T \sim 0.1$ s. This allows the system to enter a metastable state dominated by noncontractible pairs, which then live for a time set by thermal activation, $e^{\Delta/T} \sim 1$ year. The crossover to the hydrodynamic regime for $t \gg t_T$ occurs only for sufficiently large system sizes (namely, if the system can access sufficiently low monopole densities), which correspond to samples of linear dimension much larger than 10 nm

nisms at play, but we note that the precise functional form or even the asymptotic power law decay of the free monopole density implied by the model are not a requirement in order to observe a noncontractible plateau in the thermodynamic limit. Indeed, at the mean field level, *any* decay of $\rho_f(t)$ faster than $1/t$ will give rise to a plateau in the density of monopoles forming noncontractible pairs. Even if $\rho_f(t)$ does exhibit a crossover to $1/t$ behaviour at long times, the plateau will still be present in the thermodynamic limit, but will only exist for a finite period of time before it starts to decay.

4.7 Conclusions

Using a combination of Monte Carlo simulations and detailed mean field modelling, we investigated the origin of the metastable plateau that is observed in thermal quenches to low temperatures in classical spin ice [302]. Our results show that it is a consequence of the long-range nature of the Coulomb-like interactions between monopoles, combined with the system entering a non-hydrodynamic regime which is controlled by nonuniversal lattice physics. The claim that such a plateau may have been observed in recent experiments [303] therefore provides further compelling evidence for the long-range nature of the interactions between the emergent monopoles in these systems.

In particular, we have shown that when the interactions between the monopoles are truncated to finite range, the plateau reduces to a finite size effect. This is because the free monopoles in the system perform independent random walks (when their density is sufficiently low) leading to a $1/t$ decay of their density with time t . Although this is sufficient to create the “population inversion” (in which noncontractible pairs become the dominant species in the total monopole density), the slow decay of free monopoles implies that monopole-assisted decay remains effective and continues to remove noncontractible pairs from the plateau indefinitely. On the contrary, in the presence of long-range Coulomb interactions between monopoles, there exists an energetic bias in their motion across the system. At sufficiently low temperatures, which are relevant for the formation of a thermodynamic noncontractible plateau, the system enters a non-hydrodynamic regime in which the monopoles move at terminal velocity in the direction of the local force acting on them. This combination of long-range interactions and non-hydrodynamic behaviour leads to a rapid decay of the free monopole density, faster than $1/t$ and violating the law of formal kinetics. The decay of free monopoles is then sufficiently rapid to stop the monopole-assisted decay of noncontractible pairs at long times, and therefore one observes a plateau of finite density in the thermodynamic limit.

In this chapter we studied the case of sudden quenches to the target temperature. Spin ice systems and materials are well-known to exhibit long relaxation timescales at low temperatures

and a relevant and interesting question would be to investigate how much of the phenomenology observed in the present work survives in the case of ramps, where the temperature is lowered continuously to its target value, a question that is indeed of experimental importance. It would be particularly interesting to see if there is a threshold in the ramp speed beyond which the behaviour changes qualitatively. We note however that such studies, which are beyond the scope of the present work, will likely require accessing significantly lower monopole densities and therefore simulating larger system sizes, possibly beyond the current numerical capability.

Our numerical results are in quantitatively good agreement with analytics from mean field modelling. This may come as a surprise if one thinks that fluctuations in the charge density ought to bring about corrections that are not captured by mean field theory. However, emergent charges in spin ice systems are subject to a hard-core, hyperuniform constraint in their spatial distribution: The charges are born out of the underlying spins and one can easily verify that the maximal net charge that can be accumulated in a volume ℓ^3 scales as ℓ^2 (as opposed to free charge systems, where the latter can scale as ℓ^3). As a result, long-wavelength fluctuations are suppressed, and one can expect mean field calculations to be in fact rather accurate in describing spin ice behaviour. We note that the charge model introduced in this chapter is not in general subject to the same constraint. However, we impose the same initial conditions as in the spin ice system, which are therefore hyperuniform. The good agreement with mean field theory suggests that this seeding is sufficient to maintain hyperuniformity throughout the time evolution following the quench (at least within the system sizes and time scales accessible in our simulations).

Given the importance of including exchange interactions between spins beyond nearest-neighbour separation in describing the equilibrium (and out-of-equilibrium) properties of spin ice [322–325], it is pertinent to ask what the effect of such farther-ranged interactions might be on the thermal quenches discussed in the present work. Consider the inclusion of second- and third-neighbour interactions J_2 and J_3 . The latter is subdivided into J_{3a} and J_{3b} . In the special case $J_2 = -3J_{3a}$ and $J_{3b} = 0$, these interactions can be summed to give exactly the truncated (nearest-neighbour) interactions between charges: $\propto J_2 \sum_{\langle ab \rangle} Q_a Q_b$ (in addition to a shift of J_{eff}). The inclusion of such farther-ranged interactions hence modifies the short-distance physics of monopoles and leads, for example, to a modification of the barrier to activated decay of noncontractible pairs. When the interactions do not satisfy this condition, we expect nonetheless that the behaviour of the system will remain qualitatively similar provided that the long-range bias for monopole motion across the system is active during the transient (terminal velocity) regime in which the plateau is established.

Direct observation of the behaviour studied in this work requires experimental probes that measure the monopole density in spin ice materials. One could envisage using the width of the pinch points in the neutron scattering structure factor [85, 320] (with a caveat on the contribution from nearest-neighbour pairs, such as the noncontractible pairs, as discussed in Ref. [326]). Alternatively, small quenches in the magnetic field, and a measurement of the magnetisation $M(t)$ that ensues,

give access to the time evolution of the free monopole density, since $dM/dt \propto \rho_f$ [327]. Further experimental probes of monopole density in spin ice would be very much welcome in this respect.

The potential departure of long-range interacting lattice systems from a hydrodynamic description, and thence from the law of formal kinetics, is somewhat expected: At sufficiently low temperatures, the change in energy incurred by a microscopic discrete update in the system becomes larger than the thermal energy. However, one generally expects this phenomenon to affect only the short-time dynamics, and that at long times the universal hydrodynamic behaviour is recovered. Thermal quenches in spin ice demonstrate that, while this expectation must ultimately be satisfied, the altered nonuniversal, transient dynamics during times $t \lesssim 1$ s can induce very long-lived metastable states that change the behaviour of the system over a large range of ‘intermediate’ times spanning many orders of magnitude (easily growing to be of the order of 1 year or longer for experimentally relevant parameters and temperatures).

This phenomenon may play a role in other aspects of the behaviour of spin ice models and materials at low temperature (for example, a departure from hydrodynamic behaviour could be a contributing factor to the deviation from the so-called ‘quasiparticle kinetics’ in Ref. [303]). It may also be relevant to other long-range interacting natural and artificial lattice systems of interest.

5

Logarithmic growth of entanglement in a two-leg ladder

5.1 Motivation

A key theme in Chapter 3 was the emergent disorder, generated by the thermal excitation of visons, that arose in the effective description of spinons. There, our interest was in the dynamics of spinons propagating through a thermally excited but static vison background. In this chapter, we study the behaviour of a related family of quasi-one-dimensional models that host an extensive number of \mathbb{Z}_2 valued conserved charges, analogous to the visons. In contrast to Chapter 3, we are primarily interested in the case that the system is well isolated from its environment. In this context, the notion of temperature – inherently an equilibrium concept – is replaced by finite energy density, and a pure state takes the place of the canonical density matrix. This allows the system to coexist in different vison configurations simultaneously, since a generic low entanglement pure state typically has a finite projection onto (exponentially) many vison sectors. In the following, we will explore the impact of the visons on the nonequilibrium dynamics of the quasi-1D models.

The far-from-equilibrium dynamics of isolated many-body quantum systems has been a very active topic of research in multiple fields of contemporary physics, ranging from decoherence in quantum information theory to the black hole information paradox [25, 328, 329]. A central topic in this field has been the phenomenon of “many-body localisation” (MBL), by which an isolated quantum system fails to reach a local equilibrium state starting from generic initial conditions [27, 174, 330, 331]. In systems subject to strong quenched randomness, the existence of MBL can be proven under minimal assumptions [161]. Whether MBL can happen in systems with (discrete) translation invariance is a relatively subtle question [185, 190, 332–335]: in fully generic systems of this kind, it seems likely that strict MBL (i.e., a regime where a system *never* approaches equilibrium) is impossible [186, 200], at least in the conventional thermodynamic limit [336]. However, in many specific (albeit fine-tuned) models, disorder-free localisation can be established; near these fine-tuned limits, one expects the phenomenon to persist to long times, though perhaps

not asymptotically [193–199, 337].

Experimental studies of MBL have hitherto been conducted mostly on cold-atom systems and other forms of synthetic quantum matter [187, 338–343] (apart from a few studies on disordered semiconductors and superconductors [344–347], and a very recent study on phonons [348]). The key condition for disorder-free localisation—namely the presence of local conserved charges that generate intrinsic randomness at finite temperature—can also be satisfied in strongly correlated electronic systems. However, studies of disorder-free localisation in this setting have so far focused on somewhat fine-tuned models that are of limited experimental relevance, and on operators that are diagonal in the conserved charges.

Here we study specific spin ladder models that are relevant to the description of transition metal oxides [349], with an emphasis on quantities that can be measured in experiment, such as the dynamical structure factor. The models under consideration may be mapped to free fermions coupled to emergent disorder provided by local \mathbb{Z}_2 conserved charges. In contrast to previous studies, we are primarily interested in the behaviour of operators or quantities that modify the emergent disorder realisation. Such sector-changing operators are unique to systems in which the disorder is emergent, and thus the phenomenology that we consider goes beyond that of systems where the disorder is quenched. Specifically, we explore the growth of entanglement and the dynamical response of these models by relating them to Loschmidt echoes in free-fermion systems [197]. These free-fermion methods give us access to much larger system sizes than are usual in the study of MBL. The main result of this chapter is that both the entanglement dynamics and the experimentally relevant response properties of these models follow the predictions for generic *many-body* localisation: entanglement grows logarithmically in time [170, 171, 175–178] and certain dynamical correlation functions decay with anomalous power laws [350–354]. Given that the model is essentially noninteracting, this behaviour is surprising. Beyond being experimentally relevant in the study of strongly correlated materials [349], our models afford us a level of analytical understanding that allows us to elucidate why disorder-free *single particle* localisation due to emergent randomness can give rise to the same phenomenology as MBL.

We focus our attention on the square lattice compass model [111, 349, 355], which may be viewed as a quasi-one-dimensional analogue of the Kitaev honeycomb model [47] introduced in Sec. 1.4.5. This model is dual to the plaquette Ising model [356, 357], which has been explored as a prototypical model with “fracton”-like excitations, i.e., excitations whose motion is confined to reduced dimensions [358]. The relation between fractons and disorder-free MBL also remains largely unexplored in the literature (but see Ref. [359]).

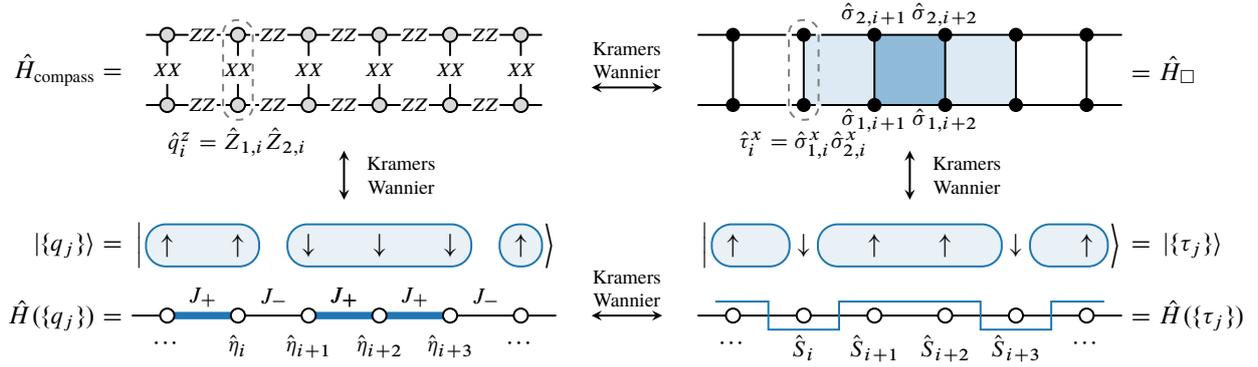


Figure 5.1: Schematic depiction of the model and its mapping to a disordered transverse field Ising model. A Kramers–Wannier duality of the compass model (5.2.1) along the rungs isolates the conserved charges $\hat{q}_i^z = \hat{Z}_{1,i} \hat{Z}_{2,i}$. Within each charge sector $\{q_j\}$ the Hamiltonian of the $\hat{\eta}$ spins $\hat{H}(\{q_j\})$ corresponds to an Ising model with nearest neighbour coupling $J_{i,i+1} = \Gamma_1 + \Gamma_2 q_i q_{i+1}$. The compass model (5.2.1) is also dual to the plaquette-Ising model (5.2.3) via a standard Kramers-Wannier transformation.

5.2 Models and mappings

We begin by introducing the compass model on a two-leg ladder [349], as illustrated in Fig. 5.1:

$$\hat{H}_{\text{compass}} = -\Delta \sum_{j=1}^L \hat{X}_{1,j} \hat{X}_{2,j} - \sum_{j=1}^{L-1} \sum_{\alpha=1}^2 \Gamma_{\alpha} \hat{Z}_{\alpha,j} \hat{Z}_{\alpha,j+1}, \quad (5.2.1)$$

where $(\hat{X}_{\alpha,j}, \hat{Z}_{\alpha,j})$ are the usual Pauli matrices on leg $\alpha = 1, 2$ and rung $j = 1, \dots, L$. Introducing the operators $\hat{q}_j^z = \hat{Z}_{1,j} \hat{Z}_{2,j}$ on each rung, $[\hat{H}, \hat{q}_j^z] = 0$ since the operators \hat{q}_j^z and $\hat{X}_{1,j} \hat{X}_{2,j}$ share either zero or two sites. This leads to an extensive number of conserved charges $\{q_j\}$, one for each rung of the ladder; since $(\hat{q}_j^z)^2 = 1$, the conserved c-numbers are $q_j = \pm 1$. The conserved charges \hat{q}_j^z are analogous to the \mathbb{Z}_2 gauge field in the Kitaev model [47] and its ladder generalisations [123, 124], introduced in Secs. 1.4.5 and 1.4.5, respectively. The presence of such local conserved charges is the hallmark of disorder-free localisation [193–199, 337].

We may then perform a 2-site version of the Kramers–Wannier duality along the rungs of the ladder to dual spin-1/2 degrees of freedom $\hat{\eta}_j$ and \hat{q}_j : $\hat{X}_{1,j} \hat{X}_{2,j} \rightarrow \hat{\eta}_j^x$, $\hat{Z}_{1,j} \rightarrow \hat{\eta}_j^z$, and $\hat{Z}_{1,j} \hat{Z}_{2,j} \rightarrow \hat{q}_j^z$. In this language, the Hamiltonian (5.2.1) becomes

$$\hat{H}_{\text{Ising}} = -\Delta \sum_{j=1}^L \hat{\eta}_j^x - \sum_{j=1}^{L-1} (\Gamma_1 + \Gamma_2 \hat{q}_j^z \hat{q}_{j+1}^z) \hat{\eta}_j^z \hat{\eta}_{j+1}^z. \quad (5.2.2)$$

There are three further equivalences to keep in mind. First, the transverse field Ising model (TFIM) (5.2.2) can be transformed, via a standard (leg-direction) Kramers–Wannier duality, to one

in which the transverse field and interaction terms are interchanged. Second, either Ising model can be mapped to free fermions via a Jordan–Wigner transformation. Third, one can undo the (rung-direction) Kramers–Wannier duality to arrive at a plaquette-Ising model with the Hamiltonian

$$\hat{H}_{\square} = -\Delta \sum_j \hat{\sigma}_{1,j}^z \hat{\sigma}_{2,j}^z \hat{\sigma}_{1,j+1}^z \hat{\sigma}_{2,j+1}^z - \sum_{j,\alpha} \Gamma_{\alpha} \hat{\sigma}_{\alpha,j}^x. \quad (5.2.3)$$

Note that the spins live on the *sites* of the square lattice here, not on the links as in the \mathbb{Z}_2 lattice gauge theory in previous chapters. Instead, it is analogous to the Xu–Moore model [360, 361] introduced in the context of $p + ip$ superconducting arrays [360–362] (which, incidentally, also appears in a dual description of the toric code in a transverse magnetic field [66]). For a discussion of boundary effects, see Appendix E. In this chapter we will treat the disorder-free spin models (5.2.1), (5.2.3) as fundamental (for the purpose of identifying local physical observables). The full set of equivalent models is captured by Fig. 5.1.

5.2.1 Exact solution

Working with tensor product states of the form $|\{q_j\}\rangle \otimes |\Phi\rangle$, the Hamiltonian (5.2.2) becomes block diagonal. Within each block, the Hamiltonian is equivalent to a standard transverse field Ising model, where the transverse magnetic field Δ is uniform, while the nearest neighbour interactions take two values $J_j \in \{J_+, J_-\}$, with $J_{\pm} \equiv \Gamma_1 \pm \Gamma_2$, depending on the configuration of the conserved charges $\{q_j^z\}$. After performing a standard Jordan–Wigner transformation, the Hamiltonian (5.2.2) may be written in terms of fermionic degrees of freedom

$$\hat{H}_{\text{TFIM}} = - \sum_{i=1}^{L-1} J_i \left(\hat{c}_i^{\dagger} \hat{c}_{i+1} + \hat{c}_{i+1}^{\dagger} \hat{c}_i + \hat{c}_i^{\dagger} \hat{c}_{i+1}^{\dagger} + \hat{c}_{i+1} \hat{c}_i \right) - \sum_{i=1}^L h_i \left(1 - 2\hat{c}_i^{\dagger} \hat{c}_i \right). \quad (5.2.4)$$

We have written the Hamiltonian in its most general form, allowing for the possibility of disorder in the magnetic fields $\{h_i\}$ and the nearest neighbour couplings $\{J_i\}$ simultaneously, since the same analysis may then be applied directly to the Kramers–Wannier dual model (5.2.3) in which the magnetic field and exchange terms are interchanged. The Hamiltonian may equivalently be written in the Majorana basis, which will prove to be particularly convenient when calculating nonequilibrium quantities. One defines $2L$ Hermitian Majorana operators $\hat{a}_i = \hat{a}_i^{\dagger}$ in place of the L complex fermions:

$$\hat{a}_{2j} = i(\hat{c}_j^{\dagger} - \hat{c}_j), \quad \hat{a}_{2j-1} = \hat{c}_j^{\dagger} + \hat{c}_j. \quad (5.2.5)$$

With the chosen normalisation, one finds that $\{\hat{a}_i, \hat{a}_j\} = 2\delta_{ij}$ (i.e., $\hat{a}_i^2 = \mathbb{1}$). In this basis, the Hamiltonian is written as

$$\hat{H}_{\text{TFIM}} = \sum_{j=1}^{L-1} iJ_j \hat{a}_{2j} \hat{a}_{2j+1} + \sum_{j=1}^L i h_j \hat{a}_{2j-1} \hat{a}_{2j} \equiv \sum_{k=1}^{2L-1} i \mathcal{J}_k \hat{a}_k \hat{a}_{k+1}, \quad (5.2.6)$$

where, after the second equality, we have defined $\mathcal{J}_{2i} = J_i$ and $\mathcal{J}_{2i-1} = h_i$. Finally, we can write (5.2.6) in terms of a real, skew-symmetric matrix $W = -W^T$ by antisymmetrising the matrix $\mathcal{J}_{ij} = \mathcal{J}_i \delta_{i,j-1}$, where $\hat{H}_{\text{TFIM}} = \frac{i}{4} \sum_{m,n} \hat{a}_m W_{mn} \hat{a}_n$. Such real skew symmetric matrices can be block-diagonalised in terms of an orthogonal matrix $Q \in O(2L)$

$$W = Q \Sigma Q^T, \quad \text{with} \quad \Sigma = \begin{pmatrix} 0 & \epsilon_1 & & & & \\ -\epsilon_1 & 0 & & & & \\ & & 0 & \epsilon_2 & & \\ & & -\epsilon_2 & 0 & & \\ & & & & \ddots & \end{pmatrix}, \quad (5.2.7)$$

where $\epsilon_n \geq 0$, a particular instance of the Schur decomposition [363]. The eigenvalues of the matrix Σ are pure imaginary, $\pm i\epsilon_n$. This block diagonal representation allows the Hamiltonian to be written in the canonical form

$$\hat{H} = \frac{i}{2} \sum_{n=1}^L \epsilon_n \hat{b}_{2n-1} \hat{b}_{2n} = \sum_{n=1}^L \epsilon_n (\hat{\gamma}_n^\dagger \hat{\gamma}_n - \frac{1}{2}), \quad (5.2.8)$$

where the complex fermions $\hat{\gamma}_n$, which diagonalise the Hamiltonian, are related to the Majorana operators \hat{b}_n via relations analogous to Eq. (5.2.5). One may alternatively diagonalise the Hamiltonian without reference to the Majorana basis by means of a Bogoliubov transformation, as described in Appendix F.

5.2.2 Anderson localisation

The spectrum of Hamiltonian (5.2.2) can then straightforwardly be constructed for any sector of the conserved charges $\{q_j\}$ using the methods described in the previous section. For random $\{q_j\}$ (e.g., in high-temperature states), the dynamics is that of Majorana fermions with random binary hopping. The Hamiltonian (5.2.2) has an eigenstate phase transition [364–366] in a given sector of $\{q_j\}$ when

$$\langle \log |\Gamma_1 + \Gamma_2 q_j q_{j+1}| \rangle = \log |\Delta|, \quad (5.2.9)$$

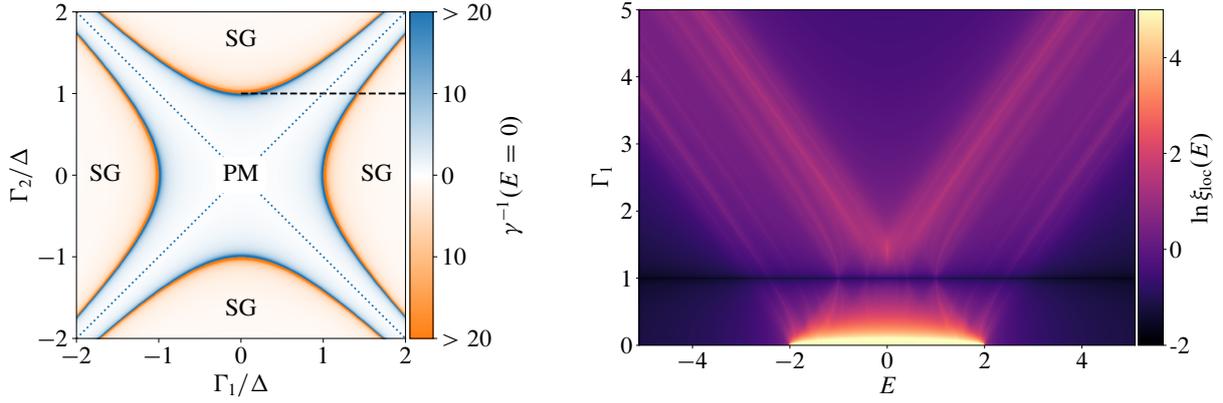


Figure 5.2: **Left:** Phase diagram of the two leg compass ladder at infinite temperature. The boundaries between the two phases, paramagnet (PM, blue) and spin glass (SG, orange), occur at $|\Gamma_1^2 - \Gamma_2^2| = \Delta^2$. This condition coincides with the vanishing of the zero-energy Lyapunov exponent, $\gamma(0)$. The blue dotted lines, $\Gamma_1 = \pm\Gamma_2$ indicate where the localisation length vanishes, and the black dashed line indicates the parameter range that we focus on. **Right:** Single particle localisation length, $\xi_{\text{loc}}(E)$, as a function of energy E along the black dashed line on the phase diagram. The system has a vanishing localisation length for equal couplings on the two legs of the ladder $\Gamma_1 = \Gamma_2$, where the chain is split into multiple disconnected sections. The localisation length was determined using standard transfer matrix techniques [145] for a ladder of length $L = 2 \times 10^5$.

where the average is over space. At infinite temperature this transition point is at $|\Gamma_1^2 - \Gamma_2^2| = \Delta^2$. It separates a random paramagnet with localised excitations—for which the order parameter autocorrelation function, $\langle \hat{\eta}_i^z(t) \hat{\eta}_i^z(0) \rangle = \langle \hat{Z}_{1,i}(t) \hat{Z}_{1,i}(0) \rangle$, vanishes—from a “spin glass” phase, in which it does not. Note that at the special value $\Gamma_1 = \Gamma_2$ the system is always paramagnetic, according to the criterion above. This follows because bonds for which $q_i q_{i+1} = -1$ are cut, and a finite segment of a system cannot undergo a phase transition. See Fig. 5.2 for the complete phase diagram. The phase transition separating these two dynamical phases is in the infinite-randomness universality class; at the transition point, the system is marginally localised with a localisation length that diverges as the single particle energy vanishes $E \rightarrow 0$ [367].

The single-particle localisation properties can be analysed more thoroughly throughout the entire spectrum using the system’s transfer matrix

$$T_i = \begin{pmatrix} -E/\Delta & -J_i/\Delta \\ 1 & 0 \end{pmatrix} \begin{pmatrix} -E/J_i & -\Delta/J_i \\ 1 & 0 \end{pmatrix} \quad (5.2.10)$$

$$= \begin{pmatrix} E^2/\Delta J_i - J_i/\Delta & E/J_i \\ -E/J_i & -\Delta/J_i \end{pmatrix}, \quad (5.2.11)$$

where the energy E parametrises the eigenvalues of the single particle Hamiltonian, defined by writing (5.2.6) in terms of complex fermions. The corresponding energy-dependent localisation length is shown in shown in Fig. 5.2, calculated using the transfer matrix methods described in

Sec. 1.5.1. Since the transfer matrix becomes diagonal at zero energy, the corresponding Lyapunov exponent $\gamma(0)$ may easily be evaluated analytically using the central limit theorem. One finds that $\gamma(0) = |\overline{\ln |J_n/\Delta|}|$, which evaluates to

$$\gamma(0) = \frac{1}{2} \text{abs ln} \frac{|\Gamma_1^2 - \Gamma_2^2|}{\Delta^2}. \quad (5.2.12)$$

As alluded to previously, $\gamma(0)$ vanishes at the phase boundaries* defined by $|\Gamma_1^2 - \Gamma_2^2| = \Delta^2$, as shown in Fig. 5.2.

As one lowers the temperature, the \hat{q}_j^z become increasingly likely to align with their neighbours, so the localisation length grows. At zero temperature, there is no randomness, and the system undergoes a ground-state phase transition that is in the Ising universality class. However, the system is localised at any finite *energy density* above the ground state.

5.3 Entanglement growth

Since the model (5.2.2) has free-fermion dynamics in any fixed sector, one can deduce that a general low-entanglement (e.g., product) initial state that is an eigenstate of all the \hat{q}_j^z will quickly saturate to area law entanglement—at least away from the critical point for that sector. If we start instead from a superposition of \hat{q}_j^z eigenstates, the entanglement exhibits unbounded slow logarithmic growth that is characteristic of MBL systems. This is the first main result of this chapter, and in what follows we explain intuitively why this happens, and then explain how one can exploit the free-fermion character of the dynamics in each sector to efficiently compute the entanglement for relatively large systems.

One can imagine “integrating out” the free fermions to arrive at an effective classical spin model with Hamiltonian $\hat{H}_{\text{eff}}(\hat{q}_j^z)$. This Hamiltonian has diagonal interactions that decay exponentially in space with the characteristic fermionic localisation length. Starting from an initial superposition, these interactions will cause slow dephasing, and thence slow entanglement growth, exactly as in Refs. [170, 171, 178]. One can illustrate this by considering a minimal example involving a 2×2 ladder. The Hamiltonian is $\hat{H}_{\text{toy}} = -(\Gamma_1 + \Gamma_2 \hat{q}_1^z \hat{q}_2^z)(\hat{c}_1^\dagger \hat{c}_2 + \hat{c}_1^\dagger \hat{c}_2^\dagger + \text{H.c.}) - \Delta \sum_j (1 - 2\hat{c}_j^\dagger \hat{c}_j)$. Considering for simplicity the sector with odd fermion parity (i.e., one fermion), the eigenstates have energies $\pm(\Gamma_1 + \Gamma_2 q_1 q_2)$. Thus, if the initial state is a superposition of different \hat{q}_j^z states, it will dephase on a time scale $\sim 1/\Gamma_2^\dagger$. The dephasing rate between pairs of \hat{q}_j falls off exponentially with distance, so at time t each \hat{q}_j is entangled with $\sim \xi \log(t/\xi)$ others [178].

*As discussed in Sec. 1.5.1, in a system of finite size, one must look instead at the fluctuations, and it can be shown that the localisation length diverges as $\sim \sqrt{L}$.

†Dephasing occurs on a different time scale in the even parity sector where the eigenstates have energies $\pm\sqrt{(\Gamma_1 + q_1 q_2 \Gamma_2)^2 + 4\Delta^2}$.



Figure 5.3: Schematic illustration of the relation between different disorder configurations in the calculation of the second Rényi entropy S_2 . The colours correspond to different disorder realisations of the charges $\{q_j\}$ in the A and B subsystems. There are four replicas of the system corresponding to two forwards (left to right) and two backwards (right to left) time evolutions. Adjacent forwards and backwards time evolutions share the same disorder configuration of either the A or B subsystem, depending on their parity, as represented by boxes of the same colour.

We now consider, more generally, an initial product state of the compass spins $|\Psi\rangle = \hat{Z}_{1,j} |\Psi\rangle = \hat{X}_{2,j} |\Psi\rangle, \forall j$. It can be written in terms of the Ising spins as

$$|\Psi\rangle = |\Phi\rangle \otimes \frac{1}{2^{L/2}} \sum_{q_j = \pm 1} |\{q_j\}\rangle, \quad (5.3.1)$$

where $\hat{\eta}_j^z |\Phi\rangle = |\Phi\rangle, \forall j$. As a result, the product state (5.3.1) has an equal-weight projection onto every charge sector.

We bipartition the system leg-wise, into two ladders A and B , each of length $L_A = L_B = L/2$. Since the reduced density matrix $\hat{\rho}_A(t)$ describes all correlations within the subsystem A , we are able to decompose it in terms of Pauli strings:

$$\hat{\rho}_A(t) = \frac{1}{2^{L_A}} \sum_{\{\mu_j\}} \text{Tr} \left[\hat{\rho}(t) \prod_{j \in A} \hat{\eta}_j^{\mu_j} \right] \prod_{j \in A} \hat{\eta}_j^{\mu_j}, \quad (5.3.2)$$

where $\mu_j = 0, 1, 2, 3$, $\hat{\eta}_j^0$ is the identity and $\hat{\eta}_j^{1,2,3} = \hat{\eta}_j^{x,y,z}$. The form of the Jordan–Wigner transformation maps the Hilbert space of the first L_A spins onto the first L_A (Majorana) fermions and thus the density matrix of the spins and of the fermions is the same [368]. Further, the transformation that maps $\hat{\eta}$ to compass spins does not mix the A and B subsystems and hence the reduced density matrices of the $\hat{\eta}$ and compass spins are unitarily equivalent.

We find that in terms of the $\hat{\eta}$ -spins

$$\text{Tr} \hat{\rho}_A^2 = \frac{1}{2^{2L}} \sum_{\{q_1\}, \{q_2\}} \text{Tr}_A \left[\text{Tr}_B \hat{U}(q_1^A, q_1^B) \hat{P}_\Phi \hat{U}^\dagger(q_2^A, q_1^B) \text{Tr}_B \hat{U}(q_2^A, q_2^B) \hat{P}_\Phi \hat{U}^\dagger(q_1^A, q_2^B) \right], \quad (5.3.3)$$

where $\hat{P}_\Phi = |\Phi\rangle \langle \Phi|$ is the projector onto the initial state of the $\hat{\eta}$ -spins, and $\hat{U}(q^A, q^B)$ is the time evolution operator with a disorder configuration specified by $\{q\} = \{q^A\} \cup \{q^B\}$. The exponentiated

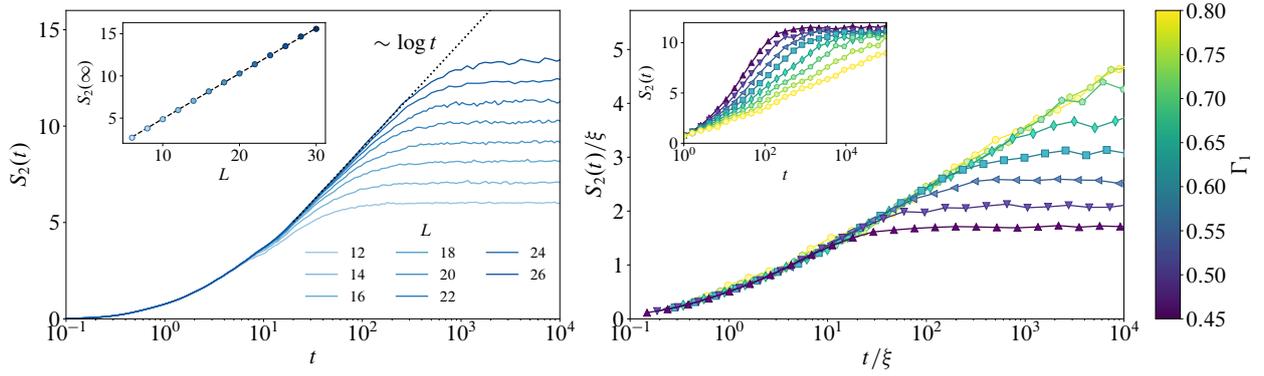


Figure 5.4: Entanglement entropy $S_2(t)$ after beginning in the translationally-invariant initial state (5.3.1) for a cut through the legs of the ladder that splits the system into two equal halves. Left panel: After some initial transient dynamics, $S_2(t)$ grows logarithmically in time, until it eventually saturates due to finite size. The saturation value is consistent with volume-law growth, as shown in the inset. Systems of size $L \leq 12$ ($N \leq 24$ spins) are calculated using exact diagonalisation, while larger system sizes are evaluated using random sampling of Eq. (5.3.3). All curves are calculated using parameters $\Gamma_2 = \Delta = 1$, and $\Gamma_1 = 1/2$. Right panel: Scaling collapse of the data for a fixed system size $L = 22$ for various values of Γ_1 , shown prior to rescaling in the inset, confirming the scaling $S_2(t) \sim \xi \log(t/\xi)$ that one may obtain within perturbation theory.

Rényi entropy $e^{-S_2(t)} \propto \text{Tr} \hat{\rho}_A^2$ may be regarded as a disorder average over two independent charge configurations $\{q_1\}$ and $\{q_2\}$. The expression includes two forwards time evolutions $\hat{U}(q^A, q^B)$, and two backwards time evolutions $\hat{U}^\dagger(q^A, q^B)$, each containing a different Hamiltonian. However, the trace enforces that the disorder configurations appearing in these Hamiltonians are not independent. For the entropy $S_\alpha(t)$ with (integer) $\alpha > 2$ there exist 2α replicas of the system with different disorder configurations correlated in the manner shown schematically in Fig. 5.3.

The expression (5.3.3) is evaluated numerically for $\alpha = 2$ using the free-fermion techniques described in the aside below. and plotted in Fig. 5.4 for $\Gamma_2 = \Delta = 1$, and $\Gamma_1 = 1/2$ (with an average localisation length $\xi \simeq 5.32$). After some initial transient dynamics, the growth of the entanglement entropy is seen to be logarithmic in time for sufficiently large systems, $S_2(t) \sim \xi \log(t/\xi)$, before finite size effects become relevant and the entropy saturates*. As shown in the inset, the late-time behaviour of S_2 is volume law: $S_2(\infty) \propto L$.

We emphasize that the logarithmic entanglement growth is a consequence of the mixing between different q -sectors in the Ising model; in a fixed q -sector, the dynamics is described by an Ising model with binary disorder, for which entanglement growth saturates (away from the critical point).

*We focus for simplicity on the second Rényi entropy, $S_2(t)$, but we expect the von Neumann entanglement entropy to exhibit similar logarithmic growth.

Aside: Free fermion techniques In free fermion systems that do not conserve particle number, only its parity, the method of Gaussian density matrix composition [368] provides a convenient way of calculating the system’s properties, both in and out of equilibrium. We make extensive use of the composition rule when calculating the entanglement entropy [Eq. (5.3.3)] and the dynamical structure factor [Eq. (5.4.1)], and will briefly review the necessary details here.

Gaussian density matrix composition. Given two normalised fermionic density matrices $\hat{\rho}_1, \hat{\rho}_2$ of the form

$$\hat{\rho}_i = \frac{1}{Z} \exp\left(\frac{1}{4} \hat{\mathbf{a}}^T W_i \hat{\mathbf{a}}\right), \quad (5.3.4)$$

where the matrices $W_i = -W_i^T$ are skew-symmetric (not necessarily Hermitian), and $\hat{\mathbf{a}}$ is a vector of Majorana operators. As shown in Ref. [368], the product $\hat{\rho}_1 \hat{\rho}_2$ is also a Gaussian density matrix. The matrix W_{12} that defines this state can be shown to satisfy $e^{W_{12}} = e^{W_1} e^{W_2}$ using the Baker–Campbell–Hausdorff (BCH) identity. However, the correlations implied by states of the form (5.3.4) are completely determined by the corresponding correlation matrix

$$\mathbb{F}_{nm} = \text{Tr}[\hat{a}_n \hat{\rho} \hat{a}_m] - \delta_{nm}. \quad (5.3.5)$$

The correlation matrix \mathbb{F} should not be confused with the parameters Γ_1 and Γ_2 that appear in the Hamiltonian. It can then be shown that, for density matrices specified by correlation matrices \mathbb{F}_1 and \mathbb{F}_2 , $\hat{\rho}[\mathbb{F}_1]$ and $\hat{\rho}[\mathbb{F}_2]$, respectively, their product satisfies the following composition rule

$$\hat{\rho}[\mathbb{F}_1] \hat{\rho}[\mathbb{F}_2] = \{\mathbb{F}_1, \mathbb{F}_2\} \hat{\rho}[\mathbb{F}_1 \times \mathbb{F}_2], \quad (5.3.6)$$

where $\{\mathbb{F}_1, \mathbb{F}_2\} \equiv \text{Tr} \hat{\rho}[\mathbb{F}_1] \hat{\rho}[\mathbb{F}_2]$, and $\mathbb{F}_1 \times \mathbb{F}_2$ is the correlation matrix of the composite density matrix. As shown in Ref. [368] the composition “ \times ” of correlation matrices is defined as

$$\mathbb{F}_1 \times \mathbb{F}_2 = \mathbb{1} - (\mathbb{1} - \mathbb{F}_2) \frac{1}{\mathbb{1} + \mathbb{F}_1 \mathbb{F}_2} (\mathbb{1} - \mathbb{F}_1). \quad (5.3.7)$$

The normalisation factor $\{\mathbb{F}_1, \mathbb{F}_2\}$ appearing in (5.3.6) may be written in terms of the spectrum of the product matrix $\mathbb{F}_1 \mathbb{F}_2$ (whose eigenvalues are doubly degenerate)

$$\{\mathbb{F}_1, \mathbb{F}_2\} = \prod_{\nu_j \in \text{Spec}(\mathbb{F}_1 \mathbb{F}_2)/2} \frac{1 + \nu_j}{2} \quad (5.3.8)$$

$$= \pm \frac{1}{2^L} \sqrt{\det |\mathbb{1} + \mathbb{F}_1 \mathbb{F}_2|}, \quad (5.3.9)$$

where the product is over half of the doubly degenerate spectrum. The unspecified sign in terms of the square root of the determinant may be resolved by writing the result in terms of Pfaffians. In particular, we find that

the normalisation factor may be written explicitly as

$$\{\Gamma_1, \Gamma_2\} = \frac{\text{Pf}(\Gamma_1^{-1} + \Gamma_2)}{2^L \text{Pf}(\Gamma_1^{-1})} = \frac{1}{(-2)^L} \text{Pf}(\Gamma_1) \text{Pf}(\Gamma_1^{-1} + \Gamma_2). \quad (5.3.10)$$

We now turn to expressing projectors $\hat{P}_\Phi = |\Phi\rangle \langle\Phi|$, appearing in a number of expressions throughout this chapter, as a Gaussian density matrix. If the state $|\Phi\rangle$ has a well-defined number of Jordan–Wigner fermions on each site in real space, then the relevant projector is

$$|\Phi\rangle \langle\Phi| = \prod_{j=1}^L \hat{P}_j^{(n_j)}, \quad (5.3.11)$$

where $\hat{P}_j^{(0)} = 1 - \hat{c}_j^\dagger \hat{c}_j$ and $\hat{P}_j^{(1)} = \hat{c}_j^\dagger \hat{c}_j$ project onto states with $n_j = 0, 1$ fermions on site j , respectively. Now, each of these projectors may be written as a Gaussian density matrix. In particular,

$$1 - \alpha \hat{c}_j^\dagger \hat{c}_j = e^{\hat{c}_j^\dagger \ln(1-\alpha) \hat{c}_j}, \quad (5.3.12)$$

where $\hat{P}_j^{(0)}$ is recovered in the limit $\alpha \rightarrow 1^-$. Conversely, for the orthogonal projector

$$\alpha^{-1} (1 + \alpha \hat{c}_j^\dagger \hat{c}_j) = \alpha^{-1} e^{\hat{c}_j^\dagger \ln(1+\alpha) \hat{c}_j}, \quad (5.3.13)$$

where now $\hat{P}_j^{(1)}$ is recovered in the limit $\alpha \rightarrow \infty$. We now proceed by writing the density matrix in terms of Majorana fermions \hat{a}_n using the relationship

$$\hat{c}_j^\dagger \hat{c}_j = \frac{1}{2} (1 + i \hat{a}_{2j-1} \hat{a}_{2j}). \quad (5.3.14)$$

Therefore, writing $\hat{\rho} = \frac{1}{Z} e^{\frac{1}{4} \sum_{mn} \hat{a}_m W_{mn} \hat{a}_n}$, the skew-symmetric matrix W decomposes into 2×2 blocks along the diagonal:

$$\ln(1 \mp \alpha) \begin{pmatrix} \hat{a}_{2\ell-1} & \hat{a}_{2\ell} \end{pmatrix} \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix} \begin{pmatrix} \hat{a}_{2\ell-1} \\ \hat{a}_{2\ell} \end{pmatrix}. \quad (5.3.15)$$

Taking the matrix hyperbolic tangent to obtain the correlation matrix, $\Gamma = \tanh(W/2)$, we arrive at

$$\Gamma_\alpha = \begin{pmatrix} 0 & i \tanh\left[\frac{1}{2} \ln(1 \mp \alpha)\right] \\ -i \tanh\left[\frac{1}{2} \ln(1 \mp \alpha)\right] & 0 \end{pmatrix} \quad (5.3.16)$$

$$\Gamma = \begin{pmatrix} 0 & \mp i \\ \pm i & 0 \end{pmatrix}^{\otimes N} = [\pm \sigma_y]^{\otimes N}, \quad (5.3.17)$$

where in the second line we have taken the appropriate limit for α . Here σ_y corresponds to the second Pauli matrix. Hence, time-dependent expressions involving the projector onto the initial state \hat{P}_Φ may be computed using the composition rule (5.3.6) and the correlation matrix (5.3.17).

Green's function approach. When the required expectation value can be written as a product of time-evolved Majorana operators, as in Eq. (5.4.2), we can use the ‘Pfaffian trick’ to map the desired correlator onto a single Pfaffian. In particular, given an ordered list of (linear combinations of) Majorana operators $\hat{\phi}_1, \hat{\phi}_2, \dots, \hat{\phi}_{2m}$, the expectation value of this list with respect to a Gaussian state Φ is given by

$$\langle \Phi | \hat{\phi}_1 \hat{\phi}_2 \cdots \hat{\phi}_{2m} | \Phi \rangle = \text{Pf}(G), \quad (5.3.18)$$

where the antisymmetric matrix G is defined by $G_{ij} = \langle \Phi | \hat{\phi}_i \hat{\phi}_j | \Phi \rangle$ for $i < j$. Applied to a time-ordered product of Majorana operators, we arrive at

$$\langle \Phi | \mathcal{T} \hat{a}_1(t_1) \hat{a}_2(t_2) \cdots \hat{a}_{2m}(t_{2m}) | \Phi \rangle = \text{Pf}(\mathcal{G}), \quad (5.3.19)$$

where $\mathcal{G}_{ij} = \langle \Phi | \mathcal{T} \hat{a}_i(t_i) \hat{a}_j(t_j) | \Phi \rangle$ for $i < j$. For $t > 0$,

$$\mathcal{G}_{ij}^> = \text{Tr}[\hat{a}_i(t) \hat{a}_j(0) | \Phi \rangle \langle \Phi |]. \quad (5.3.20)$$

Writing the time evolution of the Majoranas in terms of the unitary matrix $U(t)$, $\hat{\mathbf{a}}(t) = U(t)\hat{\mathbf{a}}(0)$, the Green's functions may be written as $\mathcal{G}_{ij}^> = [U(t)(\mathbb{1} + \mathbb{I})]_{ij}$, where \mathbb{I} is the correlation matrix of the initial state Φ .

5.4 Dynamical structure factor

Logarithmic entanglement growth, while central to the phenomenology of MBL systems, is not realistically measurable in most experiments. In what follows we consider an observable that is straightforward to measure in solid-state experiments, which we argue also exhibits signatures of MBL that are related to the logarithmic growth. Let us consider the dynamical structure factor in the basis of the compass spins $\hat{\Sigma}_{\alpha,j}$, where $\hat{\Sigma} = \hat{X}, \hat{Z}$. In particular, we are interested in the time dependence of $\langle \hat{\Sigma}_{\alpha,i}(t) \hat{\Sigma}'_{\beta,j}(0) \rangle$, where the angled brackets correspond to a finite temperature average with respect to the canonical ensemble [although analogous results are obtained using the translationally invariant initial state (5.3.1)]. The trace over charge configurations $\{q_j\}$ implies that each \hat{q}_j^μ operator that projects out of a given sector must appear an even number of times for the expectation value to be nonvanishing. As a consequence, the mixed elements XZ and ZX must vanish identically.

In the high-temperature limit, the nonzero components of the structure factor may be written as

$$\langle \hat{X}_{1,i}(t) \hat{X}_{1,j}(0) \rangle \propto \delta_{ij} \overline{\text{Tr} \left[e^{i\hat{H}(\{q\})t} e^{-i\hat{H}_i^X(\{q\}; -q_i)t} \right]}, \quad (5.4.1)$$

$$\langle \hat{Z}_{1,i}(t) \hat{Z}_{1,j}(0) \rangle \propto \delta_{ij} \overline{\text{Tr} \left[e^{i\hat{H}(\{q\})t} e^{-i\hat{H}_i^Z(\{q\})t} \right]}, \quad (5.4.2)$$

where the overline corresponds to an infinite-temperature average over the various charge sectors,

$\hat{H}_i^\mu = \hat{\eta}_i^\mu \hat{H} \hat{\eta}_i^\mu$, and $\hat{H}(\{q\}; -q_i)$ denotes that the sign of the spin q_i on site i , has been flipped with respect to the configuration $\{q\}^*$. In both cases the forwards and backwards Hamiltonians differ by some local perturbation in the real space spin basis, and may be evaluated efficiently using free-fermion techniques.

Despite the apparent similarity between the two expressions, the behaviour of the two components is markedly different. The reason for this difference is the absence (presence) of sector changing operators \hat{q}_j^x in the ZZ (XX) correlator. The ZZ correlator, being diagonal in the conserved charges, maps directly onto the order parameter correlator of the Ising Hamiltonian in Eq. (5.2.2), $\langle \hat{\eta}_i^z(t) \hat{\eta}_j^z(0) \rangle$, for which only the autocorrelation function $i = j$ is nonzero at infinite temperature [369, 370]. In the presence of emergent randomness, the behaviour of this correlator can be understood in the excited-state real-space renormalization-group (RSRG-X) framework [365]. In the paramagnetic phase, this correlator decays to zero, while in the ferromagnet it saturates to a nonzero value. (In a finite system, the correlator eventually vanishes, but on a time scale that diverges with system size.) This plateau is shown in Fig. 5.5. The ZZ correlator is therefore not sensitive to the emergent nature of the disorder, and behaves identically to a TFIM in the presence of quenched disorder. That such behaviour can occur in translationally invariant models is worthy of note, but has been observed before in a variety of contexts (see, e.g., Refs. [124, 193–199, 337]).

Conversely, the XX correlator involves both flipping Ising spins *and* changing q -sector. Since the forwards and backwards time evolutions involve *different disorder realisations*, the XX component is aware of the emergent character of the disorder. The XX correlator therefore exhibits phenomenology beyond that of conventional disordered systems, and by extension beyond that of operators that are diagonal in the local conserved charges (distinguishing our results from, e.g., Ref. [124]). The differing forwards and backwards time evolutions imply that Eq. (5.4.1) is analogous to a Loschmidt echo after a local quench. Treating the difference between the forwards and backwards time evolutions as a perturbation $\sim \epsilon(\hat{\eta}_{i-1}^z \hat{\eta}_i^z + \hat{\eta}_i^z \hat{\eta}_{i+1}^z)$ [371], we find that, in a typical q sector,

$$\langle \hat{X}_{1,i}(t) \hat{X}_{1,i}(0) \rangle \sim \prod_{n=1}^L \cos(\epsilon t [\psi_{i-1}^n \phi_i^n + \psi_i^n \phi_{i+1}^n]) \sim \left(\frac{\epsilon t}{\xi} \right)^{-c\xi}, \quad (5.4.3)$$

where the matrices ψ_j^n and ϕ_j^n diagonalise the fermionic Hamiltonian[†], and $c > 0$ is an $O(1)$ number. This correlation function is essentially the exponentiated entanglement, and represents the second main result of the chapter. We see in Fig. 5.5 that this power law decay is indeed seen in the numerics, with an exponent that is consistent with Eq. (5.4.3) (away from the critical point).

*The combined effect of commuting $\hat{\eta}_i^x$ and \hat{q}_i^x through the Hamiltonian is to change $J_\pm = \Gamma_1 \pm \Gamma_2 \rightarrow -J_\mp$ on bonds i and $i - 1$.

†Specifically, the Majorana operators that diagonalise the fermionic Hamiltonian, $\hat{H} = \frac{i}{2} \sum_n \epsilon_n \hat{\lambda}_{2n-1} \hat{\lambda}_{2n}$, are related to the real-space Majoranas via the transformations $\hat{a}_{2i} = \sum_n \psi_i^n \hat{\lambda}_{2n}$ and $\hat{a}_{2i-1} = \sum_n \phi_i^n \hat{\lambda}_{2n-1}$.

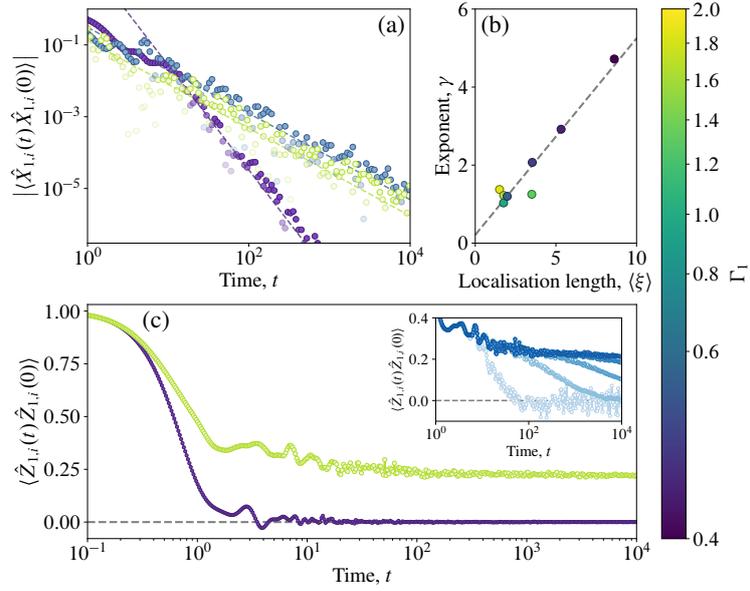


Figure 5.5: Time dependence of the diagonal elements of the infinite-temperature dynamical structure factor in the compass spins, $\hat{X}_{\alpha,j}$ and $\hat{Z}_{\alpha,j}$, for $L = 48$ ($N = 96$ spins), and $\Gamma_2 = \Delta = 1$. (a) The XX correlator exhibits a decay consistent with Eq. (5.4.3): Power law $\sim t^{-\gamma}$, with an exponent proportional to the localisation length ξ , as shown in panel (b). Conversely, the ZZ correlator (c) is diagonal in the conserved charges $\{\hat{q}_j^z\}$, and hence maps directly onto the corresponding spin correlation function of the disordered TFIM (5.2.2). The inset shows the divergence of the time scale over which the plateau decays with system size in the ferromagnetic phase (shown for $L = 8, 16, 24, 32, 40$).

5.5 Exponentially weak dephasing

In order to connect the decay of the XX component of the structure factor and the logarithmic growth of the second Rényi entropy, we can frame the latter in the language of exponentially weak dephasing, as in Eq. (5.4.3). The following arguments elucidate the correspondence between the two quantities, and provide a justification for the universal $S_2(t) \sim \xi \log(t/\xi)$ scaling observed in the numerics (Fig. 5.4). We begin by expanding the time evolution operators appearing in Eq. (5.3.3) in their eigenbases. Since the eigenstates within a typical sector at infinite temperature are exponentially localised (away from the critical point), we make the assumption that the eigenstates factorise over the two regions A and B , i.e., a given eigenstate $|\alpha\rangle$ of the composite system may be written as $|\alpha\rangle = |\alpha_A\rangle \otimes |\alpha_B\rangle$. In particular, we imagine cutting the bond between the central sites of the ladder, and take the eigenstates of each half separately. We will then reintroduce the coupling between the two halves perturbatively. The approximation of factorising the eigenstates in this way clearly breaks down at the boundary, but holds in the bulk at distances significantly greater than the localisation length ξ . It is the bulk contribution that gives rise to the logarithmic growth, and so we are not concerned with such boundary effects, which only appreciably affect the short-time

dynamics. Explicitly, each time evolution operator is written in the form

$$\hat{U}(q_A, q_B) \simeq \sum_{\alpha_A, \alpha_B} e^{-iE_{q_A, q_B}(\alpha_A, \alpha_B)t} |\alpha_A\rangle \otimes |\alpha_B\rangle \langle \alpha_A| \otimes \langle \alpha_B|. \quad (5.5.1)$$

Note that although we have assumed the eigenstates remain unchanged upon recoupling the two halves back together, we have not assumed that the same is true of the energies. In what follows we will omit the tensor product symbols for brevity. Substituting this ansatz into (5.3.3), and performing the appropriate traces over the A and B subsystems, we arrive at the expression

$$\text{Tr } \hat{\rho}_A^2 \simeq \frac{1}{2^{2L}} \sum_{\{q_1\}, \{q_2\}} \sum_{\alpha, \beta} \sum_{\gamma, \delta} e^{-iE_{11}(\alpha_A, \beta_B)t + iE_{21}(\gamma_A, \beta_B)t - iE_{22}(\gamma_A, \delta_B)t + iE_{12}(\alpha_A, \delta_B)t} \langle \alpha_A | \langle \beta_B | \hat{P}_\Phi | \gamma_A | | \beta_B \rangle \langle \gamma_A | \langle \delta_B | \hat{P}_\Phi | \alpha_A | | \delta_B \rangle, \quad (5.5.2)$$

where α_A, γ_A label the eigenstates of the A subsystem in the presence of disorder configurations q_1^A and q_2^A , and β_B, δ_B label the eigenstates of the B subsystem for configurations q_1^B and q_2^B , respectively. The notation E_{ab} is shorthand for the energy $E_{q_a^A, q_b^B}$, i.e., in the presence of disorder configuration specified by $\{q_a^A\} \cup \{q_b^B\}$. Since the initial state dependence of Eq. (5.5.2) is very weak, it is reasonable to assume that the projectors onto the initial state behave as a random matrices in the eigenbases. This allows each of the matrix elements above to be replaced by $1/2^L$ (this step can equivalently be viewed as averaging over initial states). In this way, the above expression can be written purely in terms of the eigenenergies in a manifestly positive way

$$\text{Tr } \hat{\rho}_A^2 \simeq \frac{1}{2^{4L}} \sum_{\{q_1^A\}, \{q_2^A\}} \sum_{\alpha_A, \gamma_A} \left| \sum_{\{q_1^B\}} \sum_{\beta_B} \exp \left\{ -i[E_{11}(\alpha_A, \beta_B) - E_{21}(\gamma_A, \beta_B)]t \right\} \right|^2. \quad (5.5.3)$$

We are now able to add back the connecting bond between the two halves of the system to see how the energies of the system as a whole are affected. Within second order perturbation theory, controlled by the parameter $\epsilon \ll 1$, one arrives at the following expression for the energy of the connected system

$$E_{ab}(\{\sigma_i\}) = E_{ab}^{(0)}(\{\sigma_i\}) + \epsilon \sum_i h_{ab}^i \sigma_i + \epsilon^2 \sum_{i \neq j} J_{ab}^{ij} \sigma_i \sigma_j + \dots \quad (5.5.4)$$

where the classical Ising variables $\sigma_i = 1, -1$ parametrise the occupation numbers n_i of the eigenstates, $\sigma_i = 2n_i - 1$. The local fields h_{ab}^i and the interactions J_{ab}^{ij} can be written explicitly in terms of the wave functions ψ_j^n and ϕ_j^n and the single particle energies, as shown in Appendix G. While the index i runs over the single particle energy levels, we are effectively able to treat

it as a spatial index since the eigenstates are localised. Crucially, the first order terms cancel, and the interactions typically decay with the characteristic fermionic localisation length, $J_{ab}^{ij} \sim e^{-2|i-j|/\xi}/\xi^2$ (if the two Ising spins are situated on opposite sides of the cut). Substituting this expression into (5.5.3), we find that (up to a time-dependent phase)

$$\frac{1}{2^{L/2}} \sum_{\beta_B} \exp \left\{ -i [E_{11}(\alpha_A, \beta_B) - E_{21}(\gamma_A, \beta_B)] t \right\} \propto \prod_{j \in B} \cos \left[2\epsilon^2 t \sum_{i \in A} \left(J_{11}^{ji} \sigma_i - J_{21}^{ji} \tau_i \right) \right], \quad (5.5.5)$$

where $\{\sigma_i\}$ and $\{\tau_i\}$ parametrise the occupation numbers of the states α_A and γ_A , respectively. The correspondence with the XX component of the dynamical structure factor (5.4.3) is now clear. For a given set of occupation numbers $\{\sigma_i\}$ and $\{\tau_i\}$, the summation typically behaves as $\sum_{i \in A} \left(J_{11}^{ji} \sigma_i - J_{21}^{ji} \tau_i \right) \sim e^{-2j/\xi}/\xi$, and thence $\text{Tr} \hat{\rho}_A^2 = e^{-S_2(t)} \sim (t/\xi)^{-c\xi}$, justifying the statement that (5.4.3) is essentially the exponentiated entanglement. Taking the logarithm, we arrive at the scaling behaviour $S_2(t) \sim \xi \log(t/\xi)$ observed in the right panel of Fig. 5.4. While the arguments presented here have been formulated perturbatively, exponentially weak dephasing will take place more generally as long as the effective two-body interactions decay exponentially. That is, the higher order interaction terms between the eigenstate occupation numbers of the form

$$\sum_{n=1}^{\infty} \sum_{i, j, \{k\}} K_{i\{k\}j}^{(n)} \sigma_i \sigma_{k_1} \cdots \sigma_{k_n} \sigma_j, \quad (5.5.6)$$

which give rise to an effective two body interaction of the form

$$J_{ij}^{\text{eff}}(\{\sigma_i\}) = J_{ij} + \sum_{n=1}^{\infty} \sum_{\{k\}} K_{i\{k\}j}^{(n)} \sigma_{k_1} \cdots \sigma_{k_n}, \quad (5.5.7)$$

should decay with separation as $J_{ij}^{\text{eff}} \sim e^{-|i-j|/\xi}/\xi^2$. These expressions are analogous to those that arise in the 1-bit description of the MBL phase [27, 170].

5.6 Discussion

The central result of this chapter is that quasi-1D compass and plaquette Ising models, which arise naturally in various experimental settings [349], exhibit a form of single particle disorder-free localisation that bears many of the distinctive features of MBL. In particular, we have shown that the emergent character of the disorder – which permits superpositions of different disorder realisations, and operators that modify the disorder configuration – can lead to the unbounded logarithmic growth of entanglement and anomalous power-law decay of correlation functions in spite of the noninteracting nature of the Hamiltonian. This considerably broadens the scope of

candidate materials for studying the dynamical signatures of MBL.

We established our results in a model that was solvable using free-fermion techniques; remarkably, the slow growth of entanglement, despite being inherently an interaction effect, is present in these free-fermion models because (as we explained here) integrating out the fermions gives rise to diagonal interactions and thus exponentially slow dephasing between distinct configurations of conserved variables. (Related phenomena had previously been found in out-of-time-order correlators [197, 372].) As we argued, this slow dephasing also manifests itself in more experimentally accessible variables, such as the XX component of the dynamical structure factor. Note that, while logarithmic growth of entanglement is also seen in some other models with divergent localisation lengths [373] or strong zero modes [372], the compass model in its paramagnetic phase exhibits neither of these features. Given the close parallels between the entanglement growth here and the physics of Loschmidt echoes for free fermions, the present model raises the prospect of deriving *exact* expressions for the asymptotics of entanglement and correlation functions, via solving a Riemann–Hilbert problem [374]; this is an interesting topic for future work.

A natural question our results raise is what happens for ladders with more than two legs. These systems still have one local conserved charge per rung (i.e., the product of \hat{X} operators along the rung), which can generate emergent disorder, as in the two-leg case. They are in general strongly interacting and do not admit free-fermion solutions, and are thus beyond the scope of this work. For parameters where these models have an MBL phase, their phenomenology should resemble that studied here. However, such generic interacting models will also exhibit a delocalised thermal phase. How sector-changing operators like the XX correlator behave at the many-body delocalisation transition remains an open question worthy of future consideration.

6

Conclusions and outlook

6.1 Summary

In this thesis, we have explored the interplay between nonzero temperatures and some of the characteristic emergent phenomena that occur in spin liquids.

In Chapter 2, we studied the evolution of the entanglement content in thermal states – as quantified by the logarithmic negativity – for the toric code by means of exact calculations. Intuitively, as temperature is increased, thermal fluctuations diminish the entanglement content of the system’s state as the correlations become progressively classical. However, this does not occur in a smooth manner; above a certain temperature, the negativity vanishes abruptly, a phenomenon known as sudden death. Above this temperature, the vanishing negativity indicates that the correlations in the system can be reproduced classically. The simplicity of the toric code’s spectrum allows for a transparent physical interpretation of this result: An $O(1)$ density of the lower energy defects (visons) is required to degrade the zero-temperature entanglement between two subsystems in contact with one another. However, one type of excitation alone is not sufficient to eliminate all quantum correlations, and an $O(1)$ density of the higher energy defects (spinons) is required to cause sudden death. Consequently, if the spinons are forbidden by way of a hard energetic constraint, locking in the zero temperature loop structure, quantum correlations survive up to arbitrarily high temperatures, manifesting as a slow, algebraic decay of the negativity.

In Chapter 3, we studied the effects of nontrivial mutual statistics on the propagation of quasiparticles in gapped spin liquids at finite temperature. We studied an unconstrained \mathbb{Z}_2 lattice gauge theory in the presence of a weak perpendicular magnetic field, where ring exchange leads to the emergence of the toric code Hamiltonian, and thence topological order, in the ground state sector. However, at intermediate temperatures that are comparable to the vison gap, the quantum coherence of the ground state manifold is disrupted, and the thermal state of the system corresponds to an incoherent, statistical mixture of classical ground states containing unbroken electric loops. The magnetic field also causes the spinons to acquire a nontrivial dispersion. We showed how the spinons’ dynamics may be understood in terms of a tight binding model wherein the visons act

as sources of emergent π flux. In spite of the (discrete) translational invariance of the underlying Hamiltonian, nonzero temperatures give rise to an incoherent disorder average over the emergent \mathbb{Z}_2 gauge field. Combining the approximation of self-retracing paths with numerical simulations, we were able to characterise the transport properties of spinons, showing that the system exhibits a crossover from ballistic to (sub-)diffusive behaviour at the time taken to diffuse to the nearest vison.

In Chapter 4, we looked at the nonequilibrium dynamics of monopoles in *classical* spin ice. The system is brought out of equilibrium using a thermal quench, where the temperature is suddenly lowered from its initial value, much larger than the monopole gap, to a value comparable to or lower than the gap. When brought out of equilibrium in this way, spin ice exhibits metastable states in its dynamics, which manifest as plateaux in the density of monopoles. We established that the existence of metastable plateaux in spin ice is not guaranteed; instead, its origin is dynamical in nature. Introducing noncontractible pairs as a species of quasiparticle, we formulated the emergence of the plateaux in terms of reaction–diffusion processes. Within mean field theory, we demonstrated that a plateau emerges if free monopoles are removed from the system sufficiently quickly to leave behind a profusion of noncontractible pairs. This condition is not satisfied when the interactions between the monopoles are truncated at some finite distance, in which case the plateau reduces to a finite size effect. When the monopoles interact via a long-ranged Coulomb interaction, as is expected in dipolar spin ice, the free monopoles are energetically drawn towards one another at terminal velocity, and the plateau persists in the thermodynamic limit.

We showed in Chapter 3 that the presence of an extensive number of conserved quantities can give rise to an emergent disorder average at nonzero temperature. In Chapter 5 we scrutinised the emergent character of the disorder in an analogous quasi-one-dimensional model that hosts \mathbb{Z}_2 valued conserved charges coupled to fermionic degrees of freedom. When the disorder is self-generated, a generic local operator will not only act on the fermions, but will also modify the configuration of conserved charges. Additionally, different conserved charge arrangements can be superposed. We showed that, starting from a generic low-entanglement initial state, entanglement grows logarithmically in time, just as for many body localised systems, despite the system mapping to free fermions within each charge sector. The origin of this logarithmic growth was shown to arise from exponentially weak dephasing between different configurations of the conserved charges, which inherit an effective, exponentially decaying interaction from the localised fermionic degrees of freedom. Exponentially weak dephasing also plays a prominent role in the system’s high temperature correlation functions, where it gives rise to power law decay with a continuously varying exponent.

6.2 Outlook and future directions

One of the defining traits of spin liquids is the exotic behaviour of their excitations, which inherit their properties from the proximate $T = 0$ spin liquid state. As we have shown, and as it is well known in the literature, it is often beneficial to think about these excitations as free quasiparticles hopping on a lattice (although this lattice is sometimes unrelated to the real space lattice [273]). This is certainly the case in Kitaev’s honeycomb model (see Sec. 1.4.5) and within the unconstrained lattice gauge theory from Chapter 3, where the mapping can be established exactly or within perturbation theory, respectively. This mapping renders the problem amenable to analytical and numerical approaches, and many of the thermodynamic, spectroscopic, and transport properties can be explained within linear response theory. However, in other lattice models – notably in spin ice – the excitations are inextricably linked to the spin background from which they emerge, and as a result their behaviour can be significantly more complex. Our results in Chapter 4 are a case in point, since the spin background allows for the formation of noncontractible pairs of excitations, wherein the two oppositely charged constituents are separated by a spin whose reversal does not lead to annihilation of the pair. Similar considerations apply in the context of quantum spin ice, where the spin background leads to blocked directions* and a bimodal distribution of hopping time scales [375]. In this case, the free particle description assumes the form of a hopping problem in configuration space [376], which may be approximated by a random regular graph. However, the real space system sizes that one can access with this method are limited, and in general an efficient description of quasiparticle motion in quantum spin ice remains an open question. A more complete theoretical description would allow us to better understand recent thermal conductivity measurements performed on quantum spin ice candidate materials [377].

Another striking conclusion of Chapter 4 is the importance of the short-time, nonuniversal dynamics. Indeed, we showed that, when subjected to thermal quenches, the behaviour of spin ice at long times and low temperature is determined not by hydrodynamics, but by its *initial* dynamics, where ultraviolet physics is important. At short times, performing a microscopic discrete update, such as a spin flip, can produce a change in energy comparable to or greater than temperature. One generally expects that this behaviour will manifest as a short, transient contribution to the dynamics, followed by the onset of universal behaviour. However, we showed that the initial, nonuniversal dynamics allows the system to enter a metastable state that persists for many orders of magnitude in time. This result highlights the necessity of including ultraviolet physics if we are to provide an accurate description of experiments, both numerical and in the laboratory.

Over the last few decades, we have witnessed a vast improvement in experimental capabilities, allowing us to probe electronic systems at ultrafast time scales with unprecedented spatial resolution

*In the presence of blocked directions alone, the motion of monopoles is quantum diffusive, and the Bethe lattice mapping in Chapter 3 may be used to obtain the diffusion constant.

(see, for instance, Refs. [378–381]). In addition, there has been a proliferation of experimental platforms (so-called quantum simulators) possessing a high degree of control over their constituent degrees of freedom, which are protected to a large degree from environmental influence, leading to long coherence times. Experimental platforms including ultracold atoms [382–386], trapped ions [387–389], Rydberg atoms [390], superconducting qubits [391, 392], and others [393, 394], have profoundly altered the relationship between theory and experiment, allowing us to follow the real time dynamics of quantum many body systems with single lattice site resolution.

The implications of these recent advances for the field of spin liquids are twofold. Firstly, these new quantum simulators provide an alternative avenue to study the behaviour of quantum spin liquids with an unprecedented level of control (and allowing access to nonlocal observables, which remain invisible to local experimental probes). Very recently, there have been a number of attempts to realise topological spin liquids using superconducting circuits [395], Rydberg atoms [396], and quantum annealers [53], following earlier attempts on significantly smaller systems [397–405]. The finite coherences times and high spatial resolution of these platforms further motivates the need to focus on characterising and quantifying the behaviour of many body systems on short time scales and over short distances. Secondly, the development of new experimental techniques allows for condensed matter systems to be probed in novel ways. For example, one very promising endeavour is to use two-dimensional coherent spectroscopy (2DCS) [380, 406–408] to identify new signatures of spin liquid behaviour from a system’s *nonlinear* response properties [407, 409–411]. In contrast to linear response, 2DCS is able to distinguish between a continuum of sharp modes and broadening due to a finite lifetime, thereby offering a unique fingerprint of fractionalisation [407]. Whilst the seed work has been completed, there is still much to be learnt, in particular regarding the effects of including interactions, disorder, and nonzero temperature (see however Refs. [411–413]).

We are particularly excited about the possibility of observing our results in the commercially available D-Wave quantum annealers [414, 415], which are based on an array of superconducting qubits. These machines realise a transverse field Ising model of the form $\hat{H}_{\text{TFIM}} = \sum_{i,j} J_{ij} \hat{\sigma}_i^x \hat{\sigma}_j^x + h \sum_i \hat{\sigma}_i^z$. To make contact with the results in Chapter 3, we need to obtain a four-spin interaction of the form $\hat{A}_s = \prod_{i \in s} \hat{\sigma}_i^x$ from the two-body spin interactions that are present in \hat{H}_{TFIM} . Fortunately, this can be achieved by means of combinatorial gauge symmetry, as shown in Ref. [52]. When the transverse magnetic field is weak, the emergent vison gap is suppressed by factors of h/J , and since the machines operate at nonzero temperature (of order 10 mK [414]), the experiments are performed in approximately the temperature regime studied in Chapter 3, wherein the visons (spinons) behave incoherently (coherently). The short time behaviour of the spinon density profile will be particularly important here, not least in relation to the finite decoherence time, but also to distinguish between classical and quantum diffusion, which at sufficiently long times give rise to similar behaviour.

*The couplings J_{ij} are local and can only couple spins that are connected by an edge on the so-called Chimera graph on the D-Wave 2000Q machines [416].

An additional question that arises in this context is the role of dephasing and dissipation. In this thesis, the role of environmental coupling is present only in the system's initial conditions. The system is assumed to be weakly coupled to a thermal heat bath such that the initial state of the system corresponds to one of the standard equilibrium ensembles from statistical mechanics. Time evolution thereafter is generated by the system Hamiltonian alone. However, in order to describe quantitatively the aforementioned experiments on quantum annealers, it will be necessary to treat the problem within the framework of open quantum systems. This is another context in which the spin background from which the quasiparticles emerge plays an important role: The trail of flipped spin left behind as a quasiparticle moves influences its dephasing rate. This question was partially addressed in one dimension in Ref. [417], but its extension to higher dimensions remains unresolved.

In this thesis, we have looked at a number of phenomena that arise at finite temperature in (quantum) spin liquids. A unifying theme amongst the different chapters is the presence of a finite density of thermally excited quasiparticles, inheriting their properties from the $T = 0$ spin liquid state, which have a dramatic impact on the behaviour of the system with respect to zero temperature. We have highlighted that understanding this behaviour is important not only from a practical point of view, but also as a means to uncover new fingerprints of quantum spin liquid behaviour.

A

Diffusion on other lattices

A.1 Toric code on the kagome lattice

In this appendix we calculate explicitly the generating functions for nonreversing walks on the triangular and honeycomb lattices. Both lattices appear naturally in the context of frustrated magnetism. However, as a concrete example, consider the toric code defined on the kagome lattice, as in Fig. A.1:

$$\hat{H} = -\lambda_A \sum_{\square} \hat{A}_{\square} - \lambda_B \sum_{\triangle} \hat{B}_{\triangle}, \quad (\text{A.1.1})$$

where $\hat{A}_{\square} = \prod_{i \in \square} \hat{\sigma}_i^x$, $\hat{B}_{\triangle} = \prod_{i \in \triangle} \hat{\sigma}_i^z$, corresponding to the hexagonal (\square) and triangular (\triangle , both ‘up’ and ‘down’) plaquettes of the lattice, respectively. $\lambda_{A(B)} > 0$ are the two coupling constants of the model. The operators are all mutually commuting, $[\hat{A}_{\square}, \hat{B}_{\triangle}] = 0$, since each hexagonal plaquette shares an even number of spins with any overlapping triangular plaquette.

As in the case of the square lattice, the cases $\lambda_A \ll \lambda_B$ ($\lambda_B \ll \lambda_A$) can be generated perturbatively in the ground state sector by applying a small magnetic field in the x (z) direction to a system with $\lambda_{A(B)} = 0$. We will use the terminology that the lower-energy excitations, generated perturbatively via ring exchange, correspond to the visons. If the visons reside on the triangular plaquettes, then the spinons, which live on the hexagonal plaquettes, hop on a triangular lattice. Conversely, in the opposite limiting case, if the visons live on the hexagonal plaquettes, then the spinons hop on a hexagonal lattice.

A.2 Triangular lattice

For the triangular lattice the coordination number $z = 6$ and there are thus six possible moves at each step, enumerated by the generating variables δ and ϵ : $\delta, \epsilon, \epsilon^{-1}, \delta^{-1}, \epsilon\delta^{-1}$ and $\delta\epsilon^{-1}$. The

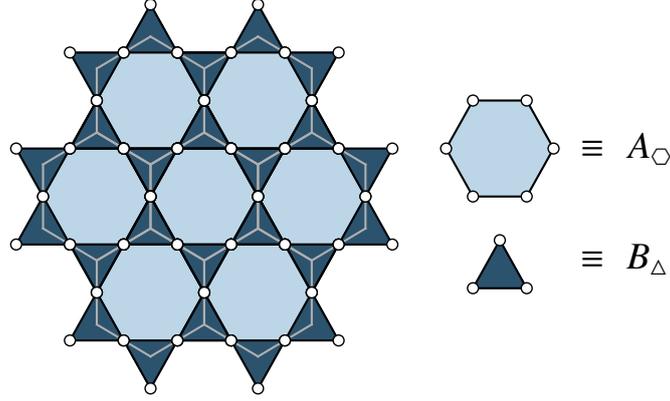


Figure A.1: A kagome lattice of spins, depicted by the white circles, and the corresponding plaquette operators, \hat{A}_{\square} and \hat{B}_{Δ} , which comprise the toric code Hamiltonian (A.1.1). The centres of the hexagonal plaquettes, \hat{A}_{\square} , form a triangular lattice, while the centres of the triangular plaquettes, \hat{B}_{Δ} , form a hexagonal lattice.

matrix which governs transitions between these various allowed moves between adjacent sites is

$$N = x \begin{pmatrix} \epsilon & \delta & \epsilon^{-1}\delta & 0 & \delta^{-1} & \epsilon\delta^{-1} \\ \epsilon & \delta & \epsilon^{-1}\delta & \epsilon^{-1} & 0 & \epsilon\delta^{-1} \\ \epsilon & \delta & \epsilon^{-1}\delta & \epsilon^{-1} & \delta^{-1} & 0 \\ 0 & \delta & \epsilon^{-1}\delta & \epsilon^{-1} & \delta^{-1} & \epsilon\delta^{-1} \\ \epsilon & 0 & \epsilon^{-1}\delta & \epsilon^{-1} & \delta^{-1} & \epsilon\delta^{-1} \\ \epsilon & \delta & 0 & \epsilon^{-1} & \delta^{-1} & \epsilon\delta^{-1} \end{pmatrix}, \quad (\text{A.2.1})$$

where the zeros enforce the nonreversing constraint imposed on the lattice walk. The initial condition

$$N_0 = x \text{diag}(\epsilon, \delta, \epsilon^{-1}\delta, \epsilon^{-1}, \delta^{-1}, \epsilon\delta^{-1}), \quad (\text{A.2.2})$$

represents the unconstrained first step. Using the general expression (3.4.15) presented in the main text, the generating function for nonreversing walks is therefore

$$\mathcal{N}(x; \delta, \epsilon) = \frac{1 - x^2}{1 - x(\epsilon + \epsilon^{-1} + \delta + \delta^{-1} + \epsilon\delta^{-1} + \delta\epsilon^{-1}) + 5x^2}. \quad (\text{A.2.3})$$

The expression for $\mathcal{R}_2(x)$ presented in the main text in (3.5.10) must be generalised to allow for the two basis vectors to be non-orthonormal, i.e., when $\mathbf{e}_i \cdot \mathbf{e}_j = a^2\delta_{ij} + a^2(1 - \delta_{ij})\cos\theta$, the expression for \mathcal{R}_2 becomes

$$\mathcal{R}_2(x) = a^2 \left\{ [(\delta\partial_{\delta})^2 + 2\cos\theta(\delta\partial_{\delta})(\epsilon\partial_{\epsilon}) + (\epsilon\partial_{\epsilon})^2] \mathcal{N} \right\} \Big|_{\delta=\epsilon=1}. \quad (\text{A.2.4})$$

By symmetry, $\delta\partial_\delta$ and $\epsilon\partial_\epsilon$ commute when acting on $\mathcal{N}(x; \delta, \epsilon)$. Evaluating the appropriate derivatives of $\mathcal{N}(x; \delta, \epsilon)$ we arrive at

$$\mathcal{R}_2(x) = \frac{4x(1+x)}{(1-5x)^2(1-x)}(2 - \cos \theta). \quad (\text{A.2.5})$$

This leads to the final expression for the diffusion constant

$$D_6 = \frac{3}{\pi} \left[3\sqrt{5} - 2 \ln \left(\frac{3 + \sqrt{5}}{3 - \sqrt{5}} \right) \right] \quad (\text{A.2.6})$$

$$= 2.72968 \dots \quad (\text{A.2.7})$$

A.3 Honeycomb lattice

The case of the honeycomb lattice ($z = 3$) is complicated slightly by its two-sublattice structure. We proceed by constructing two generating functions \mathcal{N}_{aa} and \mathcal{N}_{ab} , corresponding to walks that begin and end on the same sublattice, and walks that begin and end on complementary sublattices, respectively. Beginning with \mathcal{N}_{aa} , we divide each walk into segments of length two. Taking into account the nonreversing constraint, there are six possible transitions for each length-two segment: $\delta, \epsilon, \epsilon^{-1}, \delta^{-1}, \epsilon\delta^{-1}$ and $\delta\epsilon^{-1}$, corresponding to moves on the underlying triangular (Bravais) lattice. At each step following the initial one, *two* of these moves are disallowed by the nonreversing constraint leading to the transition matrix

$$N = x^2 \begin{pmatrix} \epsilon & \delta & 0 & 0 & \delta^{-1} & \epsilon\delta^{-1} \\ \epsilon & \delta & \epsilon^{-1}\delta & \epsilon^{-1} & 0 & 0 \\ 0 & 0 & \epsilon^{-1}\delta & \epsilon^{-1} & \delta^{-1} & \epsilon\delta^{-1} \\ \epsilon & \delta & \epsilon^{-1}\delta & \epsilon^{-1} & 0 & 0 \\ 0 & 0 & \epsilon^{-1}\delta & \epsilon^{-1} & \delta^{-1} & \epsilon\delta^{-1} \\ \epsilon & \delta & 0 & 0 & \delta^{-1} & \epsilon\delta^{-1} \end{pmatrix}, \quad (\text{A.3.1})$$

with the initial condition

$$N_0 = x^2 \text{diag}(\epsilon, \delta, \epsilon^{-1}\delta, \epsilon^{-1}, \delta^{-1}, \epsilon\delta^{-1}). \quad (\text{A.3.2})$$

These matrices lead to the generating function

$$\mathcal{N}_{aa}(x; \delta, \epsilon) = 1 + \frac{x^2 (\epsilon + \epsilon^{-1} + \delta + \delta^{-1} + \epsilon\delta^{-1} + \delta\epsilon^{-1}) - 6x^2}{1 - x^2 (\epsilon + \epsilon^{-1} + \delta + \delta^{-1} + \epsilon\delta^{-1} + \delta\epsilon^{-1} - 1) + 4x^4}. \quad (\text{A.3.3})$$

For the generating function N_{ab} , we write a walk from $a \rightarrow b$ as (i) the first step takes the walker from the a to the b sublattice, and (ii) the walker then performs a walk amongst sites belonging to the b sublattice only. This walk is implemented using the matrix

$$N = x^2 \begin{pmatrix} \epsilon & \delta & 0 & \epsilon^{-1} & 0 & \epsilon\delta^{-1} \\ \epsilon & \delta & 0 & \epsilon^{-1} & 0 & \epsilon\delta^{-1} \\ 0 & \delta & \epsilon^{-1}\delta & \epsilon^{-1} & \delta^{-1} & 0 \\ 0 & \delta & \epsilon^{-1}\delta & \epsilon^{-1} & \delta^{-1} & 0 \\ \epsilon & 0 & \epsilon^{-1}\delta & 0 & \delta^{-1} & \epsilon\delta^{-1} \\ \epsilon & 0 & \epsilon^{-1}\delta & 0 & \delta^{-1} & \epsilon\delta^{-1} \end{pmatrix}, \quad (\text{A.3.4})$$

and the initial condition

$$N_0 = xN_1 + x\epsilon^{-1}N_2 + x\delta^{-1}N_3, \quad (\text{A.3.5})$$

which represents the three possible moves in the unconstrained first step, where $N_i = \mathbf{e}_{2i} \otimes \mathbf{e}_{2i}$, and \mathbf{e}_μ ($\mu = 1, \dots, 6$) are the orthonormal basis vectors with respect to which (A.3.4) is expressed. These matrices lead to

$$\mathcal{N}_{ab}(x; \delta, \epsilon) = \frac{x(1-x^2)(1+\delta^{-1}+\epsilon^{-1})}{1-x^2(\epsilon+\epsilon^{-1}+\delta+\delta^{-1}+\epsilon\delta^{-1}+\delta\epsilon^{-1}-1)+4x^4}. \quad (\text{A.3.6})$$

The full generating function is then given by $\mathcal{N} = \mathcal{N}_{aa} + \mathcal{N}_{ab}$. However, in order to calculate \mathcal{R}_2 , one should in principle account for the fact that the b sublattice is translated by one lattice constant with respect to the a sublattice. This detail is only relevant for short times, and hence does not need to be taken into account for the calculation of the diffusion constant, which depends only on the asymptotic behaviour of $\langle \mathbf{r}^2(t) \rangle$.

Combining all of the above results gives us the generating function $\mathcal{R}_2(x)$:

$$\mathcal{R}_2(x) = \frac{2a^2}{(1-4x^2)^2(1-x^2)} \left\{ 4x^2(1+2x^2) + x(1+7x^2+4x^4) - \cos\theta [2x^2(1+2x^2) + 6x^3] \right\}, \quad (\text{A.3.7})$$

where $a = \sqrt{3}$ and $\theta = \pi/3$. Inserting these values simplifies the expression to

$$\mathcal{R}_2(x) = \frac{6x(1+2x^2)}{(1-2x)^2(1-x)(1+2x)}. \quad (\text{A.3.8})$$

Hence, we arrive at the expression

$$D_3 = \frac{3}{2\pi} \left[12\sqrt{2} - \ln \left(\frac{3 + 2\sqrt{2}}{3 - 2\sqrt{2}} \right) \right] = 3.20977\dots \quad (\text{A.3.9})$$

B | Higher order moments

Here we evaluate arbitrary moments of the density distribution in order to give a better characterisation of the spinon density profile from Chapter 3. We specialise to the case of the square lattice (i.e., $z = 4$) for convenience. The function $\mathcal{R}_{2k}(x)$ in general has a pole of order $k + 1$ at $x = (z - 1)^{-1}$. This implies that, in the long-time limit, the $2k$ th moment behaves as $\sim t^k$, consistent with the diffusive behaviour exhibited by the second moment. This is because, analogous to Eq. (3.5.12) in the main text, at long times

$$\text{Res}_{w=0} \frac{f(w)e^{itw}}{w^{k+1}} = \frac{1}{k!} (it)^k f(0) + \dots, \quad (\text{B.1})$$

if $f(w)$ is analytic at $w = 0$. The dots correspond to lower powers of t , which contribute to the transient oscillatory behaviour at short times.

We start by considering the expression for $\mathcal{R}_{2k}(x)$ derived in the main text in Eq. (3.5.10)

$$\mathcal{R}_{2k}(x) = [\nabla^{2k} \mathcal{N}](x; 1, 1) \quad (\text{B.2})$$

$$= \left[\sum_{\ell=0}^k \binom{k}{\ell} (\delta\partial_\delta)^{2\ell} (\epsilon\partial_\epsilon)^{2(k-\ell)} \mathcal{N} \right] (x; 1, 1). \quad (\text{B.3})$$

One can show that the term $(\delta\partial_\delta)^{2\ell} (\epsilon\partial_\epsilon)^{2(k-\ell)} \mathcal{N}$ gives rise to a contribution

$$\binom{k}{\ell} [2\ell!] [2(k-\ell)!] \frac{x^k (1+x)}{(1-3x)^{k+1} (1-x)^k} + \dots \quad (\text{B.4})$$

to the highest order pole $\sim (1-3x)^{-k-1}$. The dots correspond to poles of lower order that give rise to lower powers of time. Performing the summation over ℓ , we obtain

$$\sum_{\ell=0}^k \binom{k}{\ell}^2 [2\ell!] [2(k-\ell)!] = 4^k (k!)^2. \quad (\text{B.5})$$

Hence, sufficiently close to the pole at $x = 1/3$, the function \mathcal{R}_{2k} behaves as

$$\mathcal{R}_{2k}(x) \sim \frac{2^{2+k}}{3} \frac{\Gamma(k+1)^2}{(1-3x)^{k+1}}, \quad (\text{B.6})$$

where $\Gamma(x)$ is the Gamma function. The final ingredient therefore is the integral over residues, which generalises the expression (3.5.13)

$$\int_{-2\sqrt{3}}^{2\sqrt{3}} du \frac{(12-u^2)^{\frac{k+1}{2}}}{16-u^2} = \sqrt{\pi} 2^{k-2} 3^{(k+2)/2} \frac{\Gamma(\frac{3+k}{2})}{\Gamma(\frac{4+k}{2})} {}_2F_1\left(\frac{1}{2}, 1; \frac{3}{4}; \frac{3}{4}\right), \quad (\text{B.7})$$

in terms of the Gauss hypergeometric function ${}_2F_1(a, b, c; z)$. Combining the multiplicity of the highest order pole and the integral over residues, we arrive at the final exact expression for the $2k$ th moment of the density distribution in the long-time limit, for fixed k :

$$\mu_{2k} \equiv \lim_{t \rightarrow \infty} \frac{\langle \mathbf{r}^{2k}(t) \rangle}{(ht)^k} = \frac{2^{2k-1} 3^{k/2+1}}{\sqrt{\pi}} \frac{\Gamma(k+1) \Gamma(\frac{3+k}{2})}{\Gamma(\frac{4+k}{2})} {}_2F_1\left(\frac{1}{2}, 1; \frac{3}{4}; \frac{3}{4}\right). \quad (\text{B.8})$$

These moments have been checked against numerics (data not shown). As required, the special case $k = 1$ simply reduces to $2D_4$ derived in the main text. The density distribution is, however, *not* Gaussian, as evidenced by nonzero higher order cumulants. This is not an artefact of the Bethe lattice mapping, and indeed is reflected in our numerical simulations, as seen in Fig. 3.8. This feature further distinguishes the interference-driven quantum diffusion from its classical counterpart.

One can similarly compute the moments of the marginal distribution $\rho(r_x; t) = \int dr_y \rho(\mathbf{r}; t)$. One must in this case replace $\mathcal{R}_{2k}(x)$ by the function

$$\tilde{\mathcal{R}}_{2k}(x) = [(\delta \partial_\delta)^{2k} \mathcal{N}](x; 1, 1). \quad (\text{B.9})$$

A similar calculation of the residues and multiplicities (now simplified by the removal of cross terms between the generating variables δ and ϵ) gives

$$\mu_{2k}^x \equiv \lim_{t \rightarrow \infty} \frac{\langle r_x^{2k}(t) \rangle}{(ht)^k} = \frac{3^{k/2+1}}{2\sqrt{\pi}} \frac{\Gamma(2k+1) \Gamma(\frac{3+k}{2})}{\Gamma(k+1) \Gamma(\frac{4+k}{2})} {}_2F_1\left(\frac{1}{2}, 1; \frac{3}{4}; \frac{3}{4}\right). \quad (\text{B.10})$$

As one would expect, $\mu_2^x = D_z$. The first two of these exact moments were used to construct the analytical estimate of the marginal density profile in Fig. 3.8.

C | Corrections to finite size scaling exponent

Here we show how the finite size scaling exponent ν , which determines the finite size scaling behaviour of the plateau in classical spin ice with truncated interactions between the monopoles (Section 4.6.2), $\eta_p(L) \sim L^{-\nu}$, is affected by the inclusion of blocked directions.

As shown in Sec. 4.6, the expression for the exponent ν is given in terms of the ratio of the rate of monopole-assisted decay, \mathcal{R} , to the rate of monopole–antimonopole collision events, \mathcal{K} (during which the two monopoles either annihilate or form a new noncontractible pair); specifically, $\nu = 3\mathcal{R}/\mathcal{K}$. In order to estimate the ratio \mathcal{R}/\mathcal{K} microscopically, we consider a symmetrical cluster consisting of a central tetrahedron, and its first and second neighbouring tetrahedra (considering only the first nearest neighbours reproduces $\mathcal{R}/\mathcal{K} = 3/4$, i.e., blocked directions have no effect at this level). For concreteness, suppose that the central tetrahedron hosts a single, positively charged monopole. This positively charged monopole is either (1) isolated, corresponding to the calculation of the rate constant \mathcal{K} , or (2) one half of a noncontractible pair, with its negatively charged partner sitting on one of the first nearest-neighbour tetrahedra, corresponding to the calculation of \mathcal{R} . Assuming that there is an equal probability of finding the negatively charged monopole on any of the second-neighbour sites, we then deduce that the ratio of rates equals

$$\frac{\mathcal{R}}{\mathcal{K}} = \frac{P_2}{P_1}, \quad (\text{C.1})$$

where P_n ($n = 1, 2$) is the probability that a negatively charged monopole situated at second-neighbour distance is able to reach the positively charged monopole at the central site and annihilate (for $n = 2$) or either annihilate or form a new noncontractible pair (for $n = 1$).

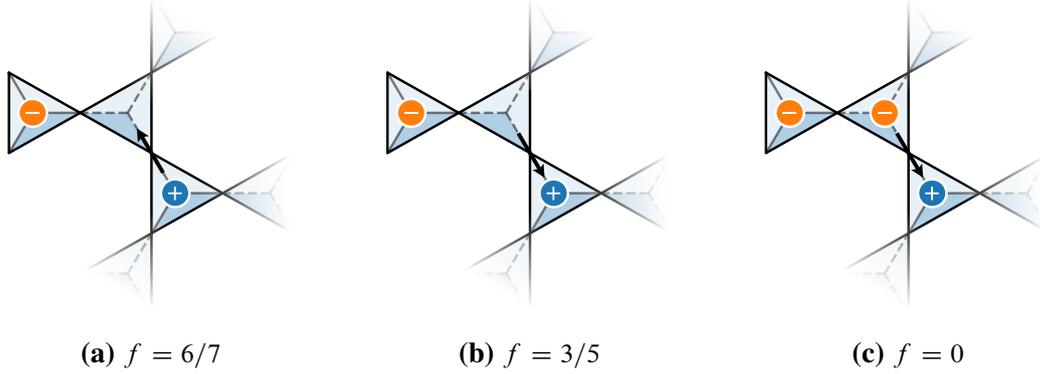


Figure C.1: Fraction f of spin configurations in which the positively and negatively charged monopoles are able to annihilate (or form a new noncontractible pair), for different configurations of spins on the central tetrahedron (which hosts the positively charged monopole).

The probabilities P_n may then be computed by enumerating the possible spin configurations:

$$P_1 = \frac{1}{N} \sum_{i=1}^N f_i = \frac{1}{4} \left[3 \times \frac{6}{7} + \frac{3}{5} \right] = \frac{111}{140} \quad (\text{C.2})$$

$$P_2 = \frac{1}{N} \sum_{i=1}^N f_i = \frac{1}{4} \left[3 \times \frac{6}{7} + 0 \right] = \frac{9}{14}, \quad (\text{C.3})$$

where $N \equiv 12$ is the number of second neighbours, and f_i is the fraction of paths starting on site i that can reach the central tetrahedron. In both expressions, the factor $6/7$ corresponds to the fraction of paths that are not blocked when the negatively charged monopole is approaching a vertex on the central tetrahedron that hosts a majority spin (Fig. C.1a). Conversely, $3/5$ of the paths are not blocked when approaching the minority spin (Fig. C.1b), unless the first-neighbour site is occupied by the negatively charged member of a noncontractible pair (Fig. C.1c), in which case the fraction is zero. Using these probabilities, we obtain the ratio of rates

$$\frac{\mathcal{R}}{\mathcal{K}} = \frac{30}{37}. \quad (\text{C.4})$$

Finally, the finite size scaling exponent is therefore $\nu = 90/37$.

D

Double charge contribution

In this appendix we show that the presence of double monopoles does not significantly alter the conclusions of our mean field modelling in Sec. 4.6 of the main text. In particular, we show by explicitly solving the mean field equations governing the density of monopoles subject to truncated interactions in the presence of double charges that, although the value of the plateau (in a finite system) is altered, the finite size scaling exponent ν remains unchanged. We argue that this feature is true more generally—further modifications of the mean field equations may change the short-time dynamics of the free monopole density, but leave its asymptotic decay ($\propto 1/t$) unchanged. This implies that the exponents derived in Sec. 4.6 are in some sense universal, while the precise value of the plateau is not (by universal we mean that the exponents are independent of how precisely the system is prepared, and are robust to the addition of terms in the mean field equations that lead to modifications of the short-time dynamics).

In addition to the species considered in Sec. 4.6, we introduce two new densities, $d_q(t)$ (where $q = \pm$), which equal the fraction of sites that host a charge $Q = \pm 2$, respectively. Notice that a double charge can always decay by reacting with any of its neighbouring tetrahedra (be them empty, occupied by a single or by a double charge), with the only exception being when it neighbours a single charge of the same sign, in which case flipping the intervening spin merely swaps the single and double charge without annihilating either of them. In principle the time evolution of the double charges depends therefore on the evolution of the single monopole density. Indeed, the average number of bonds surrounding an isolated double charge $2q$ along which it is able to decay is $4(1 - \rho_q)$ at the mean field level, i.e., assuming that each site is independent. The *asymptotic* decay of the double monopole density is however determined by neighbouring double charges of opposite sign since the number of bonds along which the pair may decay is $7/2 - 3(\rho_q + \rho_{\bar{q}})/2$ per site. Therefore, for all but the shortest times where the effect of nonzero ρ_q cannot be neglected, we expect the double charge density to decouple from the other monopole densities and to decay exponentially with a rate constant $\mathcal{K}_d \simeq 7/2$, i.e.,

$$\frac{dd_q}{dt} = -\mathcal{K}_d d_q. \quad (\text{D.1})$$

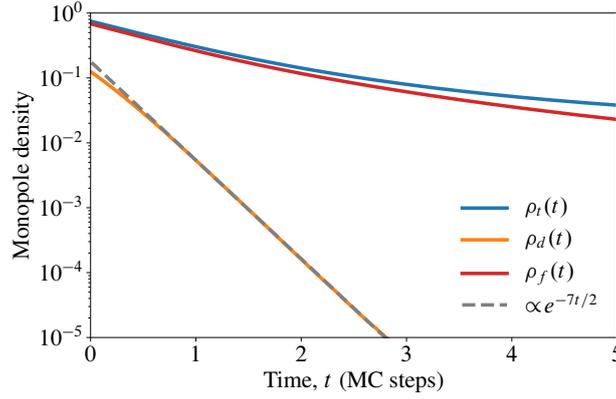


Figure D.1: Decay of the various monopole densities for a thermal quench from infinite temperature down to $T = 0.06$ K in spin ice (system size $L = 20$, i.e., 128 000 spins). The double charge density $\rho_d(t)$ is consistent with exponential decay in time with rate constant $\mathcal{K}_d = 7/2$. At very short times, $t \lesssim 1$, the effect of a nonzero free monopole density cannot be neglected, and the rate of double charge decay is reduced due to obstructed decay channels.

Adding the two equations for $q = \pm$, we obtain $\rho_d(t) = \rho_d^0 e^{-\mathcal{K}_d t}$. This expectation is indeed confirmed by our Monte Carlo simulations of CSI, where we observe asymptotic exponential decay of the total double charge density $\rho_d(t) = d_+(t) + d_-(t)$ with time (see Fig. D.1), consistent with the prediction $\mathcal{K}_d = 7/2$.

The equation governing the free charge density ρ_q must also be modified to include the effect of double monopole decay:

$$\frac{d\rho_q}{dt} = -\mathcal{K}\rho_+\rho_- + 2\mathcal{K}'_d d_q(t). \quad (\text{D.2})$$

The rate constant \mathcal{K}'_d corresponds to the spontaneous decay channel into adjacent empty sites only, implying that $\mathcal{K}'_d < \mathcal{K}_d$. Hence, the effect of including a nonzero density of double charges on the free monopole density is to add an exponentially decaying source term that corresponds to the production of free monopoles when double charges decay spontaneously. If we took into account spatial fluctuations, then we would also need to include a term $\propto (d_q \rho_{\bar{q}} - d_{\bar{q}} \rho_q)$ in this equation, but at the mean field level, charge neutrality of the single and double charges *separately* implies perfect cancellation of such a term. That is, when a single free charge q meets a double charge $2\bar{q}$, a free charge q is removed and a free charge \bar{q} is created. However, the rate at which this process occurs is identical for $q = \pm$. Substituting the exponential decay of $\rho_d(t)$ into this equation, we must solve the nonlinear equation

$$\frac{d\rho_q}{dt} + \mathcal{K}\rho_q^2 = \mathcal{K}'_d \rho_d^0 e^{-\mathcal{K}_d t}, \quad (\text{D.3})$$

for $\rho_q(t)$, in which we have made use of charge neutrality, $\rho_+(t) = \rho_-(t)$. This equation has the

exact solution

$$\rho_q(t) = y \frac{\mathcal{K}_d K_1(y) - c I_1(y)}{2\mathcal{K} K_0(y) + c I_0(y)}, \quad (\text{D.4})$$

where we have written, for convenience of notation, $y(t) \equiv 2\sqrt{\mathcal{K}\mathcal{K}'_d\rho_d^0/\mathcal{K}_d^2}e^{-\mathcal{K}_d t/2}$. The constant c is determined by the initial conditions $\rho_q(0) = \rho_q^0$, and $I_n(x)$ and $K_n(x)$ are modified Bessel functions of the first and second kind, respectively.

Finally, the expression for $\eta(t)$ must also be modified for direct comparison with our numerical results. When two double charges (of opposite sign) are adjacent to one another, the bond necessarily hosts one contractible pair and one noncontractible pair. The number of adjacent doubly occupied sites is simply proportional to $\rho_d(t)$ at long times, and the corresponding contribution to $\eta(t)$ contributes towards the kink in the noncontractible pair density observed in our numerical simulations at the characteristic decay time $t \sim \mathcal{K}_d^{-1}$ of the double charges. At later times, the equation for $\eta(t)$ remains unchanged*

$$\frac{d\eta}{dt} = -\frac{1}{2}\mathcal{R}(\rho_+ + \rho_-)\eta + \mathcal{K}'\rho_+\rho_-. \quad (\text{D.5})$$

The form of the solution is

$$\eta(t) = e^{-\mathcal{R}\Theta(t)} \left[\eta(0) - \frac{\mathcal{K}'}{\mathcal{K}} \int_0^t dt' e^{\mathcal{R}\Theta(t')} \dot{\rho}_q(t') \right], \quad (\text{D.6})$$

where we recall that $\Theta(t) \equiv \int_0^t dt' \rho_q(t')$. Hence, the asymptotic behaviour of $\eta(t)$ is directly determined by the asymptotic behaviour of $\rho_q(t)$. In order to derive this behaviour, we require the expansions of $I_n(x)$ and $K_n(x)$ for small values of the argument x [418]:

$$I_0(x) = 1 + O(x^2), \quad (\text{D.7})$$

$$I_1(x) = \frac{1}{2}x + O(x^3), \quad (\text{D.8})$$

$$K_0(x) = -\ln \frac{e^\gamma}{2}x + O(x^2 \ln x), \quad (\text{D.9})$$

$$K_1(x) = \frac{1}{x} + \frac{1}{2}x \ln x + O(x), \quad (\text{D.10})$$

*Including terms that correspond to the decay of neighbouring double charges into free monopoles gives rise to an exponentially decaying contribution to $\eta(t)$.

where $\gamma \simeq 0.5772$ is the Euler–Mascheroni constant. These expansions allow us to deduce that

$$\rho_q(t) = \frac{\mathcal{K}_d}{2\mathcal{K}} y \frac{1/y + (y/2) \ln y + O(y)}{\ln 2 - \ln e^\gamma y + c + O(y^2 \ln y)} \quad (\text{D.11})$$

$$= \frac{1}{\mathcal{K}t} + O(t^{-2}), \quad (\text{D.12})$$

independent of the initial conditions and independent of the initial rapid decay of double monopoles. The subleading term $\propto 1/t^2$ depends on the short-time dynamics through $\log y_0$ and through c . Correspondingly, for sufficiently large times,

$$\eta(t) \propto \frac{1}{(\mathcal{K}\rho_q^0 t)^{\mathcal{R}/\mathcal{K}}}. \quad (\text{D.13})$$

The constant of proportionality is slightly renormalised in the presence of double charges since the asymptotic expansion of the second term in (D.6) depends on $\int_0^\infty dt e^{\mathcal{R}\Theta} \dot{\rho}_q$, which in turn depends on the full time-dependence of $\rho_q(t)$, including its short-time dynamics. However, the *exponent* ν is insensitive to such details [being determined by the exponents of the leading terms in (D.12) and (D.13)], and the scaling arguments presented in the main text remain robust to the addition of doubly occupied sites. That is, the precise value of the plateau is sensitive to the addition of double monopoles into the model, but the finite size scaling exponent $\nu = 3\mathcal{R}/\mathcal{K}$ remains unchanged.

Similarly, when the charges are subject to mutual Coulombic interactions, if the leading term in the asymptotic expansion of $\rho_q(t)$ remains proportional to $t^{-\beta-1}$, then the leading, time-independent term in $\Theta(t) = \text{const.} + O(t^{-(1-\beta)/\beta})$ will be sensitive to the presence of double charges. Therefore, since this term contributes to the value of the plateau in the thermodynamic limit, η_∞ from (4.6.18) will be modified slightly in the presence of doubly occupied sites. However, the subleading contribution ($\sim t^{-(1-\beta)/\beta}$), which determines the finite size scaling exponent ν , will again be robust to the addition of doubly occupied sites, and the relation $\nu = 3(1-\beta)$, which relates the asymptotic decay of ρ_q to the finite size scaling behaviour, also remains unchanged.

More generally, adding further terms to our mean field equations (which depend on higher powers of the various densities) will indeed modify the short-time dynamics of $\rho_q(t)$. The precise density at which the plateau occurs in a system of finite size in the case of truncated interactions, and the value of the plateau in the thermodynamic limit in the case of long-range interactions depend – through (D.6) – on the *full* history of $\rho_q(t)$, and therefore will be modified. However, the *asymptotic* behaviour of $\rho_q(t)$, which directly determines the finite size scaling exponent ν for both types of interaction, is insensitive to such details.

E | Compass model dualities

In this appendix we describe in further detail the dualities that the compass Hamiltonian, studied in Chapter 5, possesses. As exploited in the main text, the original compass model can be transformed into a bond-disordered transverse field Ising model through the transformation (a two site version of the standard Kramers–Wannier duality)

$$\hat{q}_j^z = \hat{Z}_{1,j} \hat{Z}_{2,j} \quad \hat{\eta}_j^z = \hat{Z}_{1,j} \quad (\text{E.1})$$

$$\hat{q}_j^x = \hat{X}_{2,j} \quad \hat{\eta}_j^x = \hat{X}_{1,j} \hat{X}_{2,j} \quad (\text{E.2})$$

$$\hat{q}_j^y = \hat{Z}_{1,j} \hat{Y}_{2,j} \quad \hat{\eta}_j^y = \hat{Y}_{1,j} \hat{X}_{2,j}. \quad (\text{E.3})$$

The new spin-1/2 degrees of freedom $\hat{\eta}_j$ and \hat{q}_j commute with one another and individually satisfy the canonical angular momentum commutation relations. In these new variables, the compass Hamiltonian [i.e., Eq. (5.2.1) in the main text] becomes

$$\hat{H} = - \sum_{j=1}^{L-1} (\Gamma_1 + \Gamma_2 \hat{q}_j^z \hat{q}_{j+1}^z) \hat{\eta}_j^z \hat{\eta}_{j+1}^z - \Delta \sum_{j=1}^L \hat{\eta}_j^x. \quad (\text{E.4})$$

In this language, the local operators \hat{q}_j^z are conserved quantities. Interchanging $1 \leftrightarrow 2$ in the mapping (E.1)–(E.3) swaps the role of Γ_1 and Γ_2 in (E.4).

Alternatively, one can perform a Kramers–Wannier duality along the two legs of the ladder of the form $\hat{Z}_{\alpha,j} \hat{Z}_{\alpha,j+1} \rightarrow \hat{\sigma}_{\alpha,j}^x$, and $\hat{X}_{\alpha,j} \rightarrow \hat{\sigma}_{\alpha,j-1}^z \hat{\sigma}_{\alpha,j}^z$. This transformation gives rise to a square plaquette Ising model in the presence of a transverse field:

$$\hat{H} = -\Delta \sum_p \prod_{i \in p} \hat{\sigma}_i^z - \sum_i \Gamma_i \hat{\sigma}_i^x \quad (\text{E.5})$$

$$= -\Delta \sum_j \hat{\sigma}_{1,j}^z \hat{\sigma}_{2,j}^z \hat{\sigma}_{1,j+1}^z \hat{\sigma}_{2,j+1}^z - \sum_{j,\alpha} \Gamma_\alpha \hat{\sigma}_{\alpha,j}^x, \quad (\text{E.6})$$

where the index i labels all the spins on both legs. The second line in the equation above uses a different labelling scheme where j indexes the rungs of the ladder, and $\alpha = 1, 2$ identifies the legs,

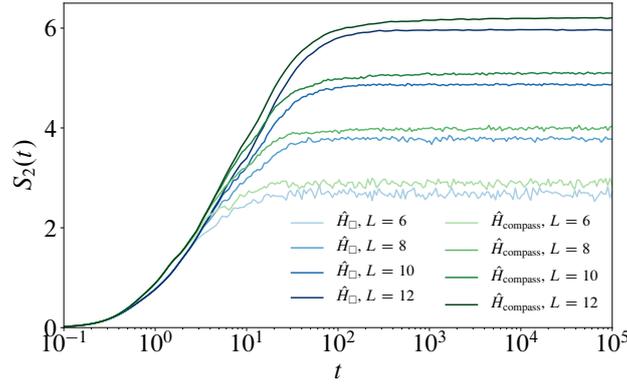


Figure E.1: Comparison of entanglement entropy for both the square plaquette (\hat{H}_{\square}) and compass (\hat{H}_{compass}) models. The two models are dual to one another via the Kramers–Wannier transformation, up to boundary effects. This leads to an $O(L^0)$ discrepancy between the two models, but does not affect the slow, logarithmic-in-time growth discussed in the main text. The curves are computed using exact diagonalisation. Taking advantage of all symmetries of the models allows us to reach $2L = 24$ spins. Parameters $\Gamma_2 = \Delta = 1$, $\Gamma_1 = 1/2$.

which are subjected to effective magnetic fields Γ_1 and Γ_2 , respectively. The conserved quantities are still products of two neighbouring spins belonging to the same rung: $\hat{\tau}_j^x = \hat{\sigma}_{1,j}^x \hat{\sigma}_{2,j}^x$, dual to plaquette operators in the original compass model, $\hat{q}_j^z \hat{q}_{j+1}^z = \hat{Z}_{1,j} \hat{Z}_{2,j} \hat{Z}_{1,j+1} \hat{Z}_{2,j+1}$. If we then perform a further Kramers–Wannier transformation along the rungs, we arrive at the Ising model, which is leg-KW dual to (E.4), i.e., where the disorder is now in the on-site magnetic field. Explicitly, implementing the transformation

$$\hat{S}^z = \hat{\sigma}_1^z \hat{\sigma}_2^z \quad \hat{\tau}^z = \hat{\sigma}_2^z \quad (\text{E.7})$$

$$\hat{S}^x = \hat{\sigma}_1^x \quad \hat{\tau}^x = \hat{\sigma}_1^x \hat{\sigma}_2^x \quad (\text{E.8})$$

$$\hat{S}^y = \hat{\sigma}_1^y \hat{\sigma}_2^z \quad \hat{\tau}^y = \hat{\sigma}_1^x \hat{\sigma}_2^y, \quad (\text{E.9})$$

we arrive at the field-disordered TFIM Hamiltonian

$$\hat{H} = -\Delta \sum_j \hat{S}_j^z \hat{S}_{j+1}^z - \sum_j (\Gamma_1 + \Gamma_2 \hat{\tau}_j^x) \hat{S}_j^x. \quad (\text{E.10})$$

If open boundary conditions are imposed on the compass spins, then this translates into fixed boundary conditions for the $\hat{\sigma}$ spin variables (and, in turn, the \hat{S} and $\hat{\tau}$ spins). The full KW transformation may be written as

$$\hat{\sigma}_{\alpha,i}^x = \hat{Z}_{\alpha,i} \hat{Z}_{\alpha,i+1} \quad (i < L), \quad \hat{\sigma}_{\alpha,L}^x = \hat{Z}_{\alpha,L} \quad (\text{E.11})$$

$$\hat{\sigma}_{\alpha,i}^z = \prod_{j \leq i} \hat{X}_{\alpha,j} \quad \forall i, \quad (\text{E.12})$$

which translates into the following Hamiltonian *including* boundary effects:

$$\hat{H} = -\Delta \sum_{j=1}^L \hat{S}_{j-1}^z \hat{S}_j^z - \sum_{j=1}^{L-1} (\Gamma_1 + \Gamma_2 \hat{t}_j^x) \hat{S}_j^x, \quad (\text{E.13})$$

where $\hat{S}_0^z = 1$. The global \mathbb{Z}_2 symmetry of the original Ising Hamiltonian, $\prod_{j=1}^L \hat{\eta}_j^x$, maps onto the conserved boundary spin \hat{S}_L^z in the dual description.

To summarise, the compass ladder is leg Kramers–Wannier dual to the square plaquette model. If open boundary conditions are imposed on the former, they manifest as fixed boundary conditions in the latter. One may equivalently impose open boundary conditions on the plaquette Ising model, leading to fixed boundary conditions imposed on the compass model. By virtue of the local duality between these models, the bulk (volume-law) contribution to the entanglement entropy is equal in the two cases, and therefore we expect to see identical unbounded logarithmic growth of entanglement in both models, up to $O(L^0)$ differences due to the boundary effects discussed above. This expectation is borne out in the numerics, as one may observe in Fig. E.1. The curves are calculated using exact diagonalisation, taking advantage of the full $\mathbb{Z}_2 \times \mathbb{Z}_2^L$ symmetry of the models. This allowed us to fully diagonalise systems of size up to and including $2L = 24$ spins (with Hilbert space dimension $\approx 1.68 \times 10^7$).

F | Bogoliubov solution to TFIM

One may also diagonalise the 1D transverse field Ising Hamiltonian encountered in Chapter 5 without reference to the Majorana basis by making use of a Bogoliubov transformation. We describe this process here. Bogoliubov fermions $\hat{\gamma}_n$ are introduced via

$$\hat{c}_i = \sum_{n=1}^L (u_{ni} \hat{\gamma}_n + v_{ni} \hat{\gamma}_n^\dagger), \quad \hat{c}_i^\dagger = \sum_{n=1}^L (u_{ni} \hat{\gamma}_n^\dagger + v_{ni} \hat{\gamma}_n). \quad (\text{F.1})$$

Since the Bogoliubov single particle Hamiltonian is real and symmetric, the matrices u_{ni} and v_{ni} may be chosen to be real. It is convenient to further introduce the quantities $\phi_{ni} = u_{ni} + v_{ni}$ and $\psi_{ni} = u_{ni} - v_{ni}$. In order to satisfy the canonical fermionic commutation relations, the matrices must be orthogonal, $\phi, \psi \in O(L)$. Explicitly,

$$\phi_{in}^T \phi_{nj} = \delta_{ij} = \psi_{in}^T \psi_{nj}, \quad (\text{F.2})$$

where repeated indices are summed over. We then proceed by writing the fermionised transverse field Ising Hamiltonian [Eq. (5.2.4) in the main text] in the form

$$\hat{H} = \sum_{i,j=1}^L \hat{c}_i^\dagger A_{ij} \hat{c}_j + \frac{1}{2} \left(\hat{c}_i^\dagger B_{ij} \hat{c}_j^\dagger - \hat{c}_i B_{ij} \hat{c}_j \right), \quad (\text{F.3})$$

where A is real and symmetric, $A = A^T$, while B is real and skew-symmetric, $B = -B^T$. In terms of the matrices A and B , the single particle Hamiltonian takes a particularly convenient form

$$\hat{H} = \frac{1}{2} \begin{pmatrix} \hat{c}^\dagger & \hat{c} \end{pmatrix} \begin{pmatrix} A & B \\ -B & -A \end{pmatrix} \begin{pmatrix} \hat{c} \\ \hat{c}^\dagger \end{pmatrix}. \quad (\text{F.4})$$

The real symmetric matrix appearing in (F.4) defines the single particle Hamiltonian $h_{ab} \in S^{2L}(\mathbb{R})$. The matrix h is unitarily (and hence spectrally) equivalent to its negative, which means that its

eigenvalues come in pairs $\pm\epsilon_n$

$$\begin{pmatrix} 0 & \mathbb{1} \\ \mathbb{1} & 0 \end{pmatrix} h \begin{pmatrix} 0 & \mathbb{1} \\ \mathbb{1} & 0 \end{pmatrix} = -h. \quad (\text{F.5})$$

Once the single particle Hamiltonian has been diagonalised, and the single particle energies ϵ_n found, the Hamiltonian can be brought into the canonical form $\hat{H} = \sum_n \epsilon_n (\hat{\gamma}_n^\dagger \hat{\gamma}_n - \frac{1}{2})$, as in Eq. (5.2.8) in the main text. In practice, one can make use of the symmetry (F.5) to diagonalise a smaller $L \times L$ matrix:

$$\sum_{j=1}^L [(A - B)(A + B)]_{ij} \phi_{nj} = \epsilon_n^2 \phi_{ni}, \quad (\text{F.6})$$

$$\sum_{j=1}^L [(A + B)(A - B)]_{ij} \psi_{nj} = \epsilon_n^2 \psi_{ni}. \quad (\text{F.7})$$

Finally, it is useful for the calculation of Lyapunov exponents to rewrite the single particle Hamiltonian by reordering the operators such that \hat{c}_i and \hat{c}_i^\dagger appear next to one another. The single particle Hamiltonian then becomes the banded matrix

$$h = \begin{pmatrix} 2h_1\tau_3 & -J_1S & & & \\ -J_1S^T & 2h_2\tau_3 & -J_2S & & \\ & -J_2S^T & 2h_3\tau_3 & \ddots & \\ & & & \ddots & \ddots \end{pmatrix}, \quad (\text{F.8})$$

where τ_3 is the third Pauli matrix, and $S = \begin{pmatrix} 1 & 1 \\ -1 & -1 \end{pmatrix}$. One again, we can make use of an expedient unitary transformation, which in this case brings the Hamiltonian into *tridiagonal* form. Specifically, we perform the local unitary transformation

$$U_j = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (\text{F.9})$$

on each site, $U = \bigoplus_{j=1}^L U_j$, which leads to the following tridiagonal matrix:

$$h = 2 \begin{pmatrix} 0 & h_1 & & & \\ h_1 & 0 & -J_2 & & \\ & -J_2 & 0 & h_2 & \\ & & h_2 & 0 & \ddots \\ & & & \ddots & \ddots \end{pmatrix}. \quad (\text{F.10})$$

G | Dephasing perturbation theory

Here we present the second order perturbation theory expressions for the effective interactions between eigenstate occupation numbers, which were used to explain the logarithmic growth of entanglement observed in Chapter 5. The perturbation theory is most easily performed in the \hat{S}_i^z language, which are the Ising spins rung Kramers–Wannier dual to the plaquette Ising model (5.2.3) (introduced in further detail in Appendix E). In this basis, the perturbation takes the form $\hat{H}' \sim \epsilon \hat{S}_i^z \hat{S}_{i+1}^z$, where the bond between sites i and $i + 1$ corresponds to the centre of the chain. In the Majorana language this perturbation may be written as $\hat{H}' = i \epsilon \hat{a}_{2i} \hat{a}_{2i+1}$. Alternatively, \hat{H}' can be written in terms of the eigenbasis of the unperturbed system, $\hat{H} = \frac{i}{2} \sum_k \epsilon_k \hat{\lambda}_{2k-1} \hat{\lambda}_{2k}$, where the two halves are independent

$$\hat{H}' = i \epsilon \sum_{k,p} \psi_{ki} \phi_{p,i+1} \hat{\lambda}_{2k} \hat{\lambda}_{2p-1}, \quad (\text{G.1})$$

or, in terms of the corresponding Bogoliubov fermions,

$$\hat{H}' = \epsilon \sum_{k,p} \psi_{ki} \phi_{p,i+1} (\hat{\gamma}_k - \hat{\gamma}_k^\dagger) (\hat{\gamma}_p + \hat{\gamma}_p^\dagger). \quad (\text{G.2})$$

Perturbation theory can now be applied using the basis $|\{n_k\}\rangle$ of unperturbed eigenstates. We are not concerned with the first order contribution, since it does not contribute to the logarithmic growth.

At second order, each of the four terms in \hat{H}' give rise to a contribution

$$\hat{\gamma}_k^\dagger \hat{\gamma}_p^\dagger : \epsilon^2 \sum_{k \neq p} \frac{-1}{-\epsilon_p - \epsilon_k} (\psi_{p,i} \phi_{k,i+1} - \psi_{k,i} \phi_{p,i+1}) \psi_{k,i} \phi_{p,i+1} (1 - n_k)(1 - n_p) \quad (\text{G.3})$$

$$\hat{\gamma}_k \hat{\gamma}_p : \epsilon^2 \sum_{k \neq p} \frac{+1}{+\epsilon_k + \epsilon_p} (\psi_{p,i} \phi_{k,i+1} - \psi_{k,i} \phi_{p,i+1}) \psi_{k,i} \phi_{p,i+1} n_k n_p \quad (\text{G.4})$$

$$\hat{\gamma}_k^\dagger \hat{\gamma}_p : \epsilon^2 \sum_{k \neq p} \frac{+1}{+\epsilon_k - \epsilon_p} (\psi_{p,i} \phi_{k,i+1} + \psi_{k,i} \phi_{p,i+1}) \psi_{k,i} \phi_{p,i+1} (1 - n_p) n_k \quad (\text{G.5})$$

$$\hat{\gamma}_k \hat{\gamma}_p^\dagger : \epsilon^2 \sum_{k \neq p} \frac{-1}{+\epsilon_p - \epsilon_k} (\psi_{p,i} \phi_{k,i+1} + \psi_{k,i} \phi_{p,i+1}) \psi_{k,i} \phi_{p,i+1} n_p (1 - n_k). \quad (\text{G.6})$$

The above terms can then be combined to arrive at the energy of a state specified by occupation numbers, $n_k = 0, 1$, of the Bogoliubov fermions

$$E(\{n_k\}) = E_0(\{n_k\}) + \epsilon E_1(\{n_k\}) + \epsilon^2 \sum_{k \neq p} [F_{kp} \delta_{n_k, n_p} + G_{kp} \delta_{n_k, 1-n_p}]. \quad (\text{G.7})$$

where the functions F_{kp} and G_{kp} may be read off from above. Alternatively, the energy can be written in terms of Ising variables $\{\sigma_k\}$ defined by $\sigma_k = 2n_k - 1$, which allows second order contribution to be written as

$$\sum_{k \neq p} [F_{k,p} \delta_{n_k, n_p} + G_{k,p} \delta_{n_k, 1-n_p}] = \frac{1}{2} \sum_{k \neq p} (F_{kp} + G_{kp}) + \frac{1}{2} \sum_{k \neq p} (F_{kp} - G_{kp}) \sigma_k \sigma_p. \quad (\text{G.8})$$

The spin-independent term can be absorbed into E_0 , and the interaction J_{kp} between the Ising spins can be defined without loss of generality as the symmetric part of $\frac{1}{2}(F_{kp} - G_{kp})$. The exponential decay of the interactions follows directly from the localised nature of the wave functions $\psi_{p,i}$ and $\phi_{p,i}$. Specifically, the (normalised) wave functions typically behave as $\psi_{p,i} \sim e^{-|i-r_p|/\xi} / \sqrt{\xi}$ (and similarly for $\phi_{p,i}$, although it should be noted that the two are not independent from one another), where r_p is the centre of localisation of the p th eigenstate. Therefore, the interactions behave qualitatively as

$$J_{kp} \sim \frac{1}{\xi^2} e^{-2|i-r_k|/\xi - 2|i-r_p|/\xi} = \frac{1}{\xi^2} e^{-2|r_k-r_p|/\xi}, \quad (\text{G.9})$$

if i is situated between the two eigenfunctions, $r_k < i < r_p$.

Bibliography

- [1] A. Tonomura, J. Endo, T. Matsuda, T. Kawasaki, and H. Ezawa, Demonstration of single-electron buildup of an interference pattern, [American Journal of Physics](#) **57**, 117–120 (1989).
- [2] E. Schrödinger, Discussion of probability relations between separated systems, [Mathematical Proceedings of the Cambridge Philosophical Society](#) **31**, 555–563 (1935).
- [3] J. S. Bell, On the Einstein Podolsky Rosen paradox, [Physics Physique Fizika](#) **1**, 195–200 (1964).
- [4] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki, Quantum entanglement, [Rev. Mod. Phys.](#) **81**, 865–942 (2009).
- [5] Alioscia Hamma, Paolo Zanardi, and Xiao-Gang Wen, String and membrane condensation on three-dimensional lattices, [Phys. Rev. B](#) **72**, 035307 (2005).
- [6] Claudio Castelnovo and Claudio Chamon, Topological order in a three-dimensional toric code at finite temperature, [Phys. Rev. B](#) **78**, 155120 (2008).
- [7] M.A. Schlosshauer, *Decoherence: And the Quantum-To-Classical Transition*, The Frontiers Collection (Springer, 2007).
- [8] Martin B. Plenio and Shashank Virmani, An introduction to entanglement measures, [Quant. Info. Comput.](#) **7**, 1–51 (2007).
- [9] M.A. Nielsen and I.L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).
- [10] Pasquale Calabrese and John Cardy, Entanglement entropy and quantum field theory, [Journal of Statistical Mechanics: Theory and Experiment](#) **2004**, P06002 (2004).
- [11] M B Hastings, An area law for one-dimensional quantum systems, [Journal of Statistical Mechanics: Theory and Experiment](#) **2007**, P08024–P08024 (2007).
- [12] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, [Rev. Mod. Phys.](#) **82**, 277–306 (2010).

- [13] Michael M. Wolf, Violation of the entropic area law for fermions, *Phys. Rev. Lett.* **96**, 010404 (2006).
- [14] Mark Fannes, Bruno Nachtergaele, and Reinhard F Werner, Finitely correlated states on quantum spin chains, *Communications in mathematical physics* **144**, 443–490 (1992).
- [15] Stellan Östlund and Stefan Rommer, Thermodynamic limit of density matrix renormalization, *Phys. Rev. Lett.* **75**, 3537–3540 (1995).
- [16] Stefan Rommer and Stellan Östlund, Class of ansatz wave functions for one-dimensional spin systems and their relation to the density matrix renormalization group, *Phys. Rev. B* **55**, 2164–2181 (1997).
- [17] Steven R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
- [18] Reinhard F. Werner, Quantum states with einstein-podolsky-rosen correlations admitting a hidden-variable model, *Phys. Rev. A* **40**, 4277–4281 (1989).
- [19] R. Horodecki, P. Horodecki, and M. Horodecki, Violating bell inequality by mixed spin- $\frac{1}{2}$ states: necessary and sufficient condition, *Physics Letters A* **200**, 340 – 344 (1995).
- [20] Juan Pablo Paz and Wojciech Hubert Zurek, Quantum limit of decoherence: Environment induced superselection of energy eigenstates, *Phys. Rev. Lett.* **82**, 5181–5185 (1999).
- [21] Alberto Frigerio, Stationary states of quantum dynamical semigroups, *Communications in Mathematical Physics* **63**, 269–276 (1978).
- [22] Herbert Spohn, Kinetic equations from hamiltonian dynamics: Markovian limits, *Rev. Mod. Phys.* **52**, 569–615 (1980).
- [23] Heinz-Peter Breuer and Francesco Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2007).
- [24] Maximilian Schlosshauer, Quantum decoherence, *Physics Reports* **831**, 1–57 (2019), quantum decoherence.
- [25] Anatoli Polkovnikov, Krishnendu Sengupta, Alessandro Silva, and Mukund Vengalattore, Colloquium: Nonequilibrium dynamics of closed interacting quantum systems, *Rev. Mod. Phys.* **83**, 863–883 (2011).
- [26] Jens Eisert, Mathis Friesdorf, and Christian Gogolin, Quantum many-body systems out of equilibrium, *Nature Physics* **11**, 124–130 (2015).

- [27] Rahul Nandkishore and David A. Huse, Many-body localization and thermalization in quantum statistical mechanics, *Annu. Rev. Condens. Matter* **6**, 15–38 (2015).
- [28] Luca D’Alessio, Yariv Kafri, Anatoli Polkovnikov, and Marcos Rigol, From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics, *Advances in Physics* **65**, 239–362 (2016).
- [29] Joshua M Deutsch, Eigenstate thermalization hypothesis, *Reports on Progress in Physics* **81**, 082001 (2018).
- [30] Mark Srednicki, Chaos and quantum thermalization, *Phys. Rev. E* **50**, 888–901 (1994).
- [31] Mark Srednicki, Thermal fluctuations in quantized chaotic systems, *Journal of Physics A: Mathematical and General* **29**, L75–L79 (1996).
- [32] Mark Srednicki, The approach to thermal equilibrium in quantized chaotic systems, *Journal of Physics A: Mathematical and General* **32**, 1163–1175 (1999).
- [33] Marcos Rigol, Vanja Dunjko, Vladimir Yurovsky, and Maxim Olshanii, Relaxation in a completely integrable many-body quantum system: An ab initio study of the dynamics of the highly excited states of 1d lattice hard-core bosons, *Phys. Rev. Lett.* **98**, 050405 (2007).
- [34] Amy C. Cassidy, Charles W. Clark, and Marcos Rigol, Generalized thermalization in an integrable lattice system, *Phys. Rev. Lett.* **106**, 140405 (2011).
- [35] Jean-Sébastien Caux and Fabian H. L. Essler, Time evolution of local observables after quenching to an integrable model, *Phys. Rev. Lett.* **110**, 257203 (2013).
- [36] Marcos Rigol, Vanja Dunjko, and Maxim Olshanii, Thermalization and its mechanism for generic isolated quantum systems, *Nature* **452**, 854–858 (2008).
- [37] Marcos Rigol, Breakdown of thermalization in finite one-dimensional systems, *Phys. Rev. Lett.* **103**, 100403 (2009).
- [38] J v Neumann, Beweis des ergodensatzes und desh-theorems in der neuen mechanik, *Zeitschrift für Physik* **57**, 30–70 (1929).
- [39] Louk Rademaker, Miguel Ortuño, and Andres M. Somoza, Many-body localization from the perspective of integrals of motion, *Annalen der Physik* **529**, 1600322 (2017).
- [40] W. Beugeling, R. Moessner, and Masudul Haque, Finite-size scaling of eigenstate thermalization, *Phys. Rev. E* **89**, 042112 (2014).

- [41] Rubem Mondaini, Keith R. Fratus, Mark Srednicki, and Marcos Rigol, Eigenstate thermalization in the two-dimensional transverse field Ising model, *Phys. Rev. E* **93**, 032104 (2016).
- [42] Maksym Serbyn, Dmitry A. Abanin, and Zlatko Papić, Quantum many-body scars and weak breaking of ergodicity, (2020), [arXiv:2011.09486 \[quant-ph\]](https://arxiv.org/abs/2011.09486).
- [43] Christopher J Turner, Alexios A Michailidis, Dmitry A Abanin, Maksym Serbyn, and Zlatko Papić, Weak ergodicity breaking from quantum many-body scars, *Nature Physics* **14**, 745–749 (2018).
- [44] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S Zibrov, Manuel Endres, Markus Greiner, *et al.*, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579–584 (2017).
- [45] Oliver Hart, Giuseppe De Tomasi, and Claudio Castelnovo, From compact localized states to many-body scars in the random quantum comb, *Phys. Rev. Research* **2**, 043267 (2020).
- [46] A.Yu. Kitaev, Fault-tolerant quantum computation by anyons, *Annals of Physics* **303**, 2–30 (2003).
- [47] Alexei Kitaev, Anyons in an exactly solved model and beyond, *Ann. Phys.* **321**, 2–111 (2006).
- [48] Giniyat Khaliullin, Orbital Order and Fluctuations in Mott Insulators, *Progress of Theoretical Physics Supplement* **160**, 155–202 (2005).
- [49] G. Jackeli and G. Khaliullin, Mott Insulators in the Strong Spin-Orbit Coupling Limit: From Heisenberg to a Quantum Compass and Kitaev Models, *Phys. Rev. Lett.* **102**, 017205 (2009).
- [50] M. Hermanns, I. Kimchi, and J. Knolle, Physics of the Kitaev Model: Fractionalization, Dynamic Correlations, and Material Connections, *Annual Review of Condensed Matter Physics* **9**, 17–33 (2018).
- [51] J. Knolle and R. Moessner, A field guide to spin liquids, *Annual Review of Condensed Matter Physics* **10**, 451–472 (2019).
- [52] Claudio Chamon, Dmitry Green, and Zhi-Cheng Yang, Constructing quantum spin liquids using combinatorial gauge symmetry, *Phys. Rev. Lett.* **125**, 067203 (2020).
- [53] Shiyu Zhou, Dmitry Green, Edward D. Dahl, and Claudio Chamon, Experimental realization of spin liquids in a programmable quantum device, (2020), [arXiv:2009.07853 \[cond-mat.str-el\]](https://arxiv.org/abs/2009.07853).

- [54] Simon Trebst, Philipp Werner, Matthias Troyer, Kirill Shtengel, and Chetan Nayak, Breakdown of a topological phase: Quantum phase transition in a loop gas model with tension, *Phys. Rev. Lett.* **98**, 070602 (2007).
- [55] Claudio Castelnovo, Simon Trebst, and Matthias Troyer, Understanding quantum phase transitions, (CRC Press, Boca Raton, 2010) Chap. Topological Order and Quantum Criticality.
- [56] Lucile Savary and Leon Balents, Quantum spin liquids: a review, *Reports on Progress in Physics* **80**, 016502 (2017).
- [57] Alexei Kitaev and Chris Laumann, Topological phases and quantum computation, (2009), [arXiv:0904.2771 \[cond-mat.mes-hall\]](https://arxiv.org/abs/0904.2771) .
- [58] Torbjörn Einarsson, Fractional statistics on a torus, *Phys. Rev. Lett.* **64**, 1995–1998 (1990).
- [59] Hong-Chen Jiang, Zhenghan Wang, and Leon Balents, Identifying topological order by entanglement entropy, *Nature Physics* **8**, 902–905 (2012).
- [60] Hong-Chen Jiang, Rajiv R. P. Singh, and Leon Balents, Accuracy of topological entanglement entropy on finite cylinders, *Phys. Rev. Lett.* **111**, 107205 (2013).
- [61] Alexei Kitaev and John Preskill, Topological entanglement entropy, *Phys. Rev. Lett.* **96**, 110404 (2006).
- [62] Michael Levin and Xiao-Gang Wen, Detecting topological order in a ground state wave function, *Phys. Rev. Lett.* **96**, 110405 (2006).
- [63] M. B. Hastings and Xiao-Gang Wen, Quasiadiabatic continuation of quantum states: The stability of topological ground-state degeneracy and emergent gauge invariance, *Phys. Rev. B* **72**, 045141 (2005).
- [64] I. S. Tupitsyn, A. Kitaev, N. V. Prokof'ev, and P. C. E. Stamp, Topological multicritical point in the phase diagram of the toric code model and three-dimensional lattice gauge higgs model, *Phys. Rev. B* **82**, 085114 (2010).
- [65] Julien Vidal, Sébastien Dusuel, and Kai Phillip Schmidt, Low-energy effective theory of the toric code model in a parallel magnetic field, *Phys. Rev. B* **79**, 033109 (2009).
- [66] Julien Vidal, Ronny Thomale, Kai Phillip Schmidt, and Sébastien Dusuel, Self-duality and bound states of the toric code model in a transverse field, *Phys. Rev. B* **80**, 081104 (2009).
- [67] Fengcheng Wu, Youjin Deng, and Nikolay Prokof'ev, Phase diagram of the toric code model in a parallel magnetic field, *Phys. Rev. B* **85**, 195104 (2012).

- [68] Franz J. Wegner, Duality in generalized ising models and phase transitions without local order parameters, *Journal of Mathematical Physics* **12**, 2259–2272 (1971).
- [69] John B. Kogut, An introduction to lattice gauge theory and spin systems, *Rev. Mod. Phys.* **51**, 659–713 (1979).
- [70] Eduardo Fradkin and Leonard Susskind, Order and disorder in gauge systems and magnets, *Phys. Rev. D* **17**, 2637–2658 (1978).
- [71] P Pfeuty and R J Elliott, The ising model with a transverse field. II. ground state properties, *Journal of Physics C: Solid State Physics* **4**, 2370–2385 (1971).
- [72] Heiko Rieger and Naoki Kawashima, Application of a continuous time cluster algorithm to the two-dimensional random quantum ising ferromagnet, *The European Physical Journal B-Condensed Matter and Complex Systems* **9**, 233–236 (1999).
- [73] Eduardo Fradkin, *Field Theories of Condensed Matter Physics*, 2nd ed. (Cambridge University Press, 2013).
- [74] W. F. Giauque and Muriel F. Ashley, Molecular Rotation in Ice at 10°K. Free Energy of Formation and Entropy of Water, *Phys. Rev.* **43**, 81–82 (1933).
- [75] W. F. Giauque and J. W. Stout, The Entropy of Water and the Third Law of Thermodynamics. The Heat Capacity of Ice from 15 to 273°K. *Journal of the American Chemical Society* **58**, 1144–1150 (1936).
- [76] Linus Pauling, The structure and entropy of ice and of other crystals with some randomness of atomic arrangement, *Journal of the American Chemical Society* **57**, 2680–2684 (1935).
- [77] J. D. Bernal and R. H. Fowler, A theory of water and ionic solution, with particular reference to hydrogen and hydroxyl ions, *The Journal of Chemical Physics* **1**, 515–548 (1933).
- [78] J. F. Nagle, Lattice Statistics of Hydrogen Bonded Crystals. I. The Residual Entropy of Ice, *Journal of Mathematical Physics* **7**, 1484–1491 (1966).
- [79] Jason S. Gardner, Michel J. P. Gingras, and John E. Greedan, Magnetic pyrochlore oxides, *Rev. Mod. Phys.* **82**, 53–107 (2010).
- [80] C. Lacroix, P. Mendels, and F. Mila, *Introduction to Frustrated Magnetism: Materials, Experiments, Theory*, Springer Series in Solid-State Sciences (Springer Berlin Heidelberg, 2011).

- [81] L. D. C. Jaubert and P. C. W. Holdsworth, Magnetic monopole dynamics in spin ice, *Journal of Physics: Condensed Matter* **23**, 164222 (2011).
- [82] P. W. Anderson, Ordering and antiferromagnetism in ferrites, *Phys. Rev.* **102**, 1008–1013 (1956).
- [83] S. V. Isakov, K. Gregor, R. Moessner, and S. L. Sondhi, Dipolar spin correlations in classical pyrochlore magnets, *Phys. Rev. Lett.* **93**, 167204 (2004).
- [84] C. Castelnovo, R. Moessner, and S.L. Sondhi, Spin ice, fractionalization, and topological order, *Annual Review of Condensed Matter Physics* **3**, 35–55 (2012).
- [85] T. Fennell, P. P. Deen, A. R. Wildes, K. Schmalzl, D. Prabhakaran, A. T. Boothroyd, R. J. Aldus, D. F. McMorrow, and S. T. Bramwell, Magnetic Coulomb Phase in the Spin Ice $\text{Ho}_2\text{Ti}_2\text{O}_7$, *Science* **326**, 415–417 (2009).
- [86] Hiroaki Kadowaki, Naohiro Doi, Yuji Aoki, Yoshikazu Tabata, Taku J Sato, Jeffrey W Lynn, Kazuyuki Matsuhira, and Zenji Hiroi, Observation of magnetic monopoles in spin ice, *Journal of the Physical Society of Japan* **78**, 103706–103706 (2009).
- [87] S. V. Isakov, R. Moessner, and S. L. Sondhi, Why spin ice obeys the ice rules, *Phys. Rev. Lett.* **95**, 217201 (2005).
- [88] Claudio Castelnovo, Roderich Moessner, and Shivaji L Sondhi, Magnetic monopoles in spin ice, *Nature* **451**, 42 (2008).
- [89] Claudio Castelnovo and Claudio Chamon, Topological order in a three-dimensional toric code at finite temperature, *Phys. Rev. B* **78**, 155120 (2008).
- [90] C. Castelnovo, R. Moessner, and S. L. Sondhi, Debye-hückel theory for spin ice at low temperature, *Phys. Rev. B* **84**, 144435 (2011).
- [91] IA Ryzhkin, Magnetic relaxation in rare-earth oxide pyrochlores, *Journal of Experimental and Theoretical Physics* **101**, 481–486 (2005).
- [92] Ludovic DC Jaubert and Peter CW Holdsworth, Signature of magnetic monopole and dirac string dynamics in spin ice, *Nature Physics* **5**, 258 (2009).
- [93] Roger G. Melko, Byron C. den Hertog, and Michel J. P. Gingras, Long-range order at low temperatures in dipolar spin ice, *Phys. Rev. Lett.* **87**, 067203 (2001).
- [94] Roger G Melko and Michel J P Gingras, Monte carlo studies of the dipolar spin ice model, *Journal of Physics: Condensed Matter* **16**, R1277–R1319 (2004).

- [95] M.J. Harris, S.T. Bramwell, T. Zeiske, D.F. McMorrow, and P.J.C. King, Magnetic structures of highly frustrated pyrochlores, *Journal of Magnetism and Magnetic Materials* **177-181**, 757–762 (1998), international Conference on Magnetism (Part II).
- [96] H. Fukazawa, R. G. Melko, R. Higashinaka, Y. Maeno, and M. J. P. Gingras, Magnetic anisotropy of the spin-ice compound $\text{Dy}_2\text{Ti}_2\text{O}_7$, *Phys. Rev. B* **65**, 054410 (2002).
- [97] Steven T. Bramwell, Michel J. P. Gingras, and Peter C. W. Holdsworth, Spin ice, in *Frustrated Spin Systems*, edited by H T Diep (WORLD SCIENTIFIC, 2013) 2nd ed., pp. 383–474.
- [98] M J P Gingras and P A McClarty, Quantum spin ice: a search for gapless quantum spin liquids in pyrochlore magnets, *Reports on Progress in Physics* **77**, 056501 (2014).
- [99] Michael Hermele, Matthew P. A. Fisher, and Leon Balents, Pyrochlore photons: The $U(1)$ spin liquid in a $S = \frac{1}{2}$ three-dimensional frustrated magnet, *Phys. Rev. B* **69**, 064404 (2004).
- [100] Kate A. Ross, Lucile Savary, Bruce D. Gaulin, and Leon Balents, Quantum excitations in quantum spin ice, *Phys. Rev. X* **1**, 021002 (2011).
- [101] A.M. Polyakov, *Gauge Fields and Strings*, Contemporary concepts in physics (Harwood Academic Publishers, 1987).
- [102] Argha Banerjee, Sergei V. Isakov, Kedar Damle, and Yong Baek Kim, Unusual liquid state of hard-core bosons on the pyrochlore lattice, *Phys. Rev. Lett.* **100**, 047208 (2008).
- [103] Nic Shannon, Olga Sikora, Frank Pollmann, Karlo Penc, and Peter Fulde, Quantum ice: A quantum monte carlo study, *Phys. Rev. Lett.* **108**, 067204 (2012).
- [104] R. Moessner and S. L. Sondhi, Three-dimensional resonating-valence-bond liquids and their excitations, *Phys. Rev. B* **68**, 184512 (2003).
- [105] Ji ř Chaloupka, George Jackeli, and Giniyat Khaliullin, Kitaev-Heisenberg Model on a Honeycomb Lattice: Possible Exotic Phases in Iridium Oxides $A_2\text{IrO}_3$, *Phys. Rev. Lett.* **105**, 027204 (2010).
- [106] Stephen M. Winter, Ying Li, Harald O. Jeschke, and Roser Valentí, Challenges in design of kitaev materials: Magnetic interactions from competing energy scales, *Phys. Rev. B* **93**, 214431 (2016).
- [107] Jeffrey G. Rau, Eric Kin-Ho Lee, and Hae-Young Kee, Generic spin model for the honeycomb iridates beyond the kitaev limit, *Phys. Rev. Lett.* **112**, 077204 (2014).

- [108] Masahiko G. Yamada, Hiroyuki Fujita, and Masaki Oshikawa, Designing kitaev spin liquids in metal-organic frameworks, *Phys. Rev. Lett.* **119**, 057202 (2017).
- [109] Andrei Catuneanu, Youhei Yamaji, Gideon Wachtel, Yong Baek Kim, and Hae-Young Kee, Path to stable quantum spin liquids in spin-orbit coupled correlated materials, *npj Quantum Materials* **3**, 1–6 (2018).
- [110] Matthias Gohlke, Gideon Wachtel, Youhei Yamaji, Frank Pollmann, and Yong Baek Kim, Quantum spin liquid signatures in Kitaev-like frustrated magnets, *Phys. Rev. B* **97**, 075126 (2018).
- [111] Zohar Nussinov and Jeroen van den Brink, Compass models: Theory and physical motivations, *Rev. Mod. Phys.* **87**, 1–59 (2015).
- [112] Patrick A. Lee, An end to the drought of quantum spin liquids, *Science* **321**, 1306–1307 (2008).
- [113] Leon Balents, Spin liquids in frustrated magnets, *Nature* **464**, 199–208 (2010).
- [114] William Witczak-Krempa, Gang Chen, Yong Baek Kim, and Leon Balents, Correlated quantum phenomena in the strong spin-orbit regime, *Annual Review of Condensed Matter Physics* **5**, 57–82 (2014).
- [115] Jeffrey G. Rau, Eric Kin-Ho Lee, and Hae-Young Kee, Spin-orbit physics giving rise to novel phases in correlated systems: Iridates and related materials, *Annual Review of Condensed Matter Physics* **7**, 195–221 (2016).
- [116] Yi Zhou, Kazushi Kanoda, and Tai-Kai Ng, Quantum spin liquid states, *Rev. Mod. Phys.* **89**, 025003 (2017).
- [117] Elliott H. Lieb, Flux phase of the half-filled band, *Phys. Rev. Lett.* **73**, 2158–2161 (1994).
- [118] G. Baskaran, Saptarshi Mandal, and R. Shankar, Exact Results for Spin Dynamics and Fractionalization in the Kitaev Model, *Phys. Rev. Lett.* **98**, 247201 (2007).
- [119] Hong Yao and Xiao-Liang Qi, Entanglement Entropy and Entanglement Spectrum of the Kitaev Model, *Phys. Rev. Lett.* **105**, 080501 (2010).
- [120] G. Kells, A. T. Bolukbasi, V. Lahtinen, J. K. Slingerland, J. K. Pachos, and J. Vala, Topological Degeneracy and Vortex Manipulation in Kitaev’s Honeycomb Model, *Phys. Rev. Lett.* **101**, 240404 (2008).

- [121] G. Kells, J. K. Slingerland, and J. Vala, Description of Kitaev’s honeycomb model with toric-code stabilizers, *Phys. Rev. B* **80**, 125415 (2009).
- [122] S Mandal, R Shankar, and G Baskaran, RVB gauge theory and the topological degeneracy in the honeycomb Kitaev model, *Journal of Physics A: Mathematical and Theoretical* **45**, 335304 (2012).
- [123] Xiao-Yong Feng, Guang-Ming Zhang, and Tao Xiang, Topological characterization of quantum phase transitions in a spin-1/2 model, *Phys. Rev. Lett.* **98**, 087204 (2007).
- [124] Alexandros Metavitsiadis and Wolfram Brenig, Thermal transport in a one-dimensional \mathbb{Z}_2 spin liquid, *Phys. Rev. B* **96**, 041115 (2017).
- [125] Han-Dong Chen and Zohar Nussinov, Exact results of the Kitaev model on a hexagonal lattice: spin states, string and brane correlators, and anyonic excitations, *Journal of Physics A: Mathematical and Theoretical* **41**, 075001 (2008).
- [126] P. W. Anderson, Absence of diffusion in certain random lattices, *Phys. Rev.* **109**, 1492–1505 (1958).
- [127] Ping Sheng, *Scattering and localization of classical waves in random media*, Vol. 8 (World Scientific, 1990).
- [128] Diederik S Wiersma, Paolo Bartolini, Ad Lagendijk, and Roberto Righini, Localization of light in a disordered medium, *Nature* **390**, 671–673 (1997).
- [129] Juliette Billy, Vincent Josse, Zhanchun Zuo, Alain Bernard, Ben Hambrecht, Pierre Lugan, David Clément, Laurent Sanchez-Palencia, Philippe Bouyer, and Alain Aspect, Direct observation of Anderson localization of matter waves in a controlled disorder, *Nature* **453**, 891–894 (2008).
- [130] Giacomo Roati, Chiara D’Errico, Leonardo Fallani, Marco Fattori, Chiara Fort, Matteo Zaccanti, Giovanni Modugno, Michele Modugno, and Massimo Inguscio, Anderson localization of a non-interacting Bose–Einstein condensate, *Nature* **453**, 895–898 (2008).
- [131] Hefei Hu, A Strybulevych, JH Page, Sergey E Skipetrov, and Bart A van Tiggelen, Localization of ultrasound in a three-dimensional elastic network, *Nature Physics* **4**, 945–948 (2008).
- [132] E. N. Economou, C. M. Soukoulis, and A. D. Zdetsis, Localized states in disordered systems as bound states in potential wells, *Phys. Rev. B* **30**, 1686–1694 (1984).

- [133] D.J. Thouless, Electrons in disordered systems and the theory of localization, *Physics Reports* **13**, 93–142 (1974).
- [134] C Castellani and L Peliti, *Journal of Physics A: Mathematical and General* **19**, L429–L432 (1986).
- [135] Ferdinand Evers and Alexander D. Mirlin, Anderson transitions, *Rev. Mod. Phys.* **80**, 1355–1417 (2008).
- [136] J T Chalker and P D Coddington, Percolation, quantum tunnelling and the integer hall effect, *Journal of Physics C: Solid State Physics* **21**, 2665–2679 (1988).
- [137] F. Evers, A. Mildenberger, and A. D. Mirlin, Multifractality of wave functions at the quantum hall transition revisited, *Phys. Rev. B* **64**, 241303 (2001).
- [138] A. De Luca, B. L. Altshuler, V. E. Kravtsov, and A. Scardicchio, Anderson localization on the bethe lattice: Nonergodicity of extended states, *Phys. Rev. Lett.* **113**, 046806 (2014).
- [139] B. L. Altshuler, E. Cuevas, L. B. Ioffe, and V. E. Kravtsov, Nonergodic phases in strongly disordered random regular graphs, *Phys. Rev. Lett.* **117**, 156601 (2016).
- [140] N. F. Mott and W. D. Twose, The theory of impurity conduction, *Advances in Physics* **10**, 107–163 (1961).
- [141] I Ya Gol'dshtein, Stanislav Alekseevich Molchanov, and Leonid Andreevich Pastur, A pure point spectrum of the stochastic one-dimensional schrödinger operator, *Functional Analysis and Its Applications* **11**, 1–8 (1977).
- [142] E. Abrahams, P. W. Anderson, D. C. Licciardello, and T. V. Ramakrishnan, Scaling theory of localization: Absence of quantum diffusion in two dimensions, *Phys. Rev. Lett.* **42**, 673–676 (1979).
- [143] V. I. Oseledec, A multiplicative ergodic theorem. Ljapunov characteristic numbers for dynamical systems, *Trudy Moskov. Mat. Obšč.* **19** (1968), English Translation. *Trans. Moscow Math. Soc.* **19** (1968), 197–231.
- [144] Madabusi S Raghunathan, A proof of oseledec's multiplicative ergodic theorem, *Israel Journal of Mathematics* **32**, 356–362 (1979).
- [145] B Kramer and A MacKinnon, Localization: theory and experiment, *Reports on Progress in Physics* **56**, 1469–1564 (1993).

- [146] C. W. J. Beenakker, Random-matrix theory of quantum transport, *Rev. Mod. Phys.* **69**, 731–808 (1997).
- [147] E.N. Economou, *Green's Functions in Quantum Physics*, Springer Series in Solid-State Sciences (Springer Berlin Heidelberg, 2013).
- [148] F.M. Izrailev, A.A. Krokhin, and N.M. Makarov, Anomalous localization in low-dimensional systems with correlated disorder, *Physics Reports* **512**, 125–254 (2012).
- [149] D. J. Thouless, *Ill-condensed matter*, (North-Holland, Amsterdam, 1979).
- [150] J L Pichard, The one-dimensional Anderson model: scaling and resonances revisited, *Journal of Physics C: Solid State Physics* **19**, 1519–1535 (1986).
- [151] Timothy A. L. Ziman, Localization and spectral singularities in random chains, *Phys. Rev. Lett.* **49**, 337–340 (1982).
- [152] Freeman J. Dyson, The dynamics of a disordered linear chain, *Phys. Rev.* **92**, 1331–1338 (1953).
- [153] George Theodorou and Morrel H. Cohen, Extended states in a one-dimensional system with off-diagonal disorder, *Phys. Rev. B* **13**, 4597–4601 (1976).
- [154] T. P. Eggarter and R. Riedinger, Singular behavior of tight-binding chains with off-diagonal disorder, *Phys. Rev. B* **18**, 569–575 (1978).
- [155] D C Herbert and R Jones, Localized states in disordered systems, *Journal of Physics C: Solid State Physics* **4**, 1145–1161 (1971).
- [156] D J Thouless, A relation between the density of states and range of localization for one dimensional random systems, *Journal of Physics C: Solid State Physics* **5**, 77–81 (1972).
- [157] H. Cheraghchi, S. M. Fazeli, and K. Esfarjani, Localization-delocalization transition in a one one-dimensional system with long-range correlated off-diagonal disorder, *Phys. Rev. B* **72**, 174207 (2005).
- [158] E. N. Economou, Determination of the conductivity in disordered systems by the potential-well analogy, *Phys. Rev. B* **31**, 7710–7713 (1985).
- [159] D C Licciardello and D J Thouless, Conductivity and mobility edges in disordered systems. II. further calculations for the square and diamond lattices, *Journal of Physics C: Solid State Physics* **11**, 925–936 (1978).

- [160] L. Fleishman and P. W. Anderson, Interactions and the anderson transition, *Phys. Rev. B* **21**, 2366–2377 (1980).
- [161] John Z. Imbrie, On many-body localization for quantum spin chains, *Journal of Statistical Physics* **163**, 998–1048 (2016).
- [162] Timothy C. Berkelbach and David R. Reichman, Conductivity of disordered quantum lattice models at infinite temperature: Many-body localization, *Phys. Rev. B* **81**, 224429 (2010).
- [163] O. S. Barišić and P. Prelovšek, Conductivity in a disordered one-dimensional system of interacting fermions, *Phys. Rev. B* **82**, 161106 (2010).
- [164] Kartiek Agarwal, Sarang Gopalakrishnan, Michael Knap, Markus Müller, and Eugene Demler, Anomalous diffusion and griffiths effects near the many-body localization transition, *Phys. Rev. Lett.* **114**, 160401 (2015).
- [165] Yevgeny Bar Lev, Guy Cohen, and David R. Reichman, Absence of diffusion in an interacting system of spinless fermions on a one-dimensional disordered lattice, *Phys. Rev. Lett.* **114**, 100601 (2015).
- [166] Robin Steinigeweg, Jacek Herbrych, Frank Pollmann, and Wolfram Brenig, Typicality approach to the optical conductivity in thermal and many-body localized phases, *Phys. Rev. B* **94**, 180401 (2016).
- [167] Cécile Monthus and Thomas Garel, Many-body localization transition in a lattice model of interacting fermions: Statistics of renormalized hoppings in configuration space, *Phys. Rev. B* **81**, 134202 (2010).
- [168] Arijeet Pal and David A. Huse, Many-body localization phase transition, *Phys. Rev. B* **82**, 174411 (2010).
- [169] David J. Luitz, Nicolas Laflorencie, and Fabien Alet, Extended slow dynamical regime close to the many-body localization transition, *Phys. Rev. B* **93**, 060201 (2016).
- [170] David A. Huse, Rahul Nandkishore, and Vadim Oganesyan, Phenomenology of fully many-body-localized systems, *Phys. Rev. B* **90**, 174202 (2014).
- [171] Maksym Serbyn, Z. Papić, and Dmitry A. Abanin, Local conservation laws and the structure of the many-body localized states, *Phys. Rev. Lett.* **111**, 127201 (2013).
- [172] Anushya Chandran, Isaac H. Kim, Guifre Vidal, and Dmitry A. Abanin, Constructing local integrals of motion in the many-body localized phase, *Phys. Rev. B* **91**, 085425 (2015).

- [173] Dmitry A. Abanin and Zlatko Papić, Recent progress in many-body localization, *Annalen der Physik* **529**, 1700169 (2017).
- [174] Dmitry A. Abanin, Ehud Altman, Immanuel Bloch, and Maksym Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [175] Marko Žnidarič, Tomaž Prosen, and Peter Prelovšek, Many-body localization in the heisenberg XXZ magnet in a random field, *Phys. Rev. B* **77**, 064426 (2008).
- [176] Jens H. Bardarson, Frank Pollmann, and Joel E. Moore, Unbounded growth of entanglement in models of many-body localization, *Phys. Rev. Lett.* **109**, 017202 (2012).
- [177] Ronen Vosk and Ehud Altman, Dynamical quantum phase transitions in random spin chains, *Phys. Rev. Lett.* **112**, 217204 (2014).
- [178] Maksym Serbyn, Z. Papić, and Dmitry A. Abanin, Universal slow growth of entanglement in interacting strongly disordered systems, *Phys. Rev. Lett.* **110**, 260601 (2013).
- [179] Don N. Page, Average entropy of a subsystem, *Phys. Rev. Lett.* **71**, 1291–1294 (1993).
- [180] A. Smith, J. Knolle, D. L. Kovrizhin, and R. Moessner, Disorder-free localization, *Phys. Rev. Lett.* **118**, 266601 (2017).
- [181] Adam Smith, *Disorder-free localization* (Springer, 2019).
- [182] Yu Kagan and LA Maksimov, Localization in a system of interacting particles diffusing in a regular crystal, *Zhurnal Eksperimental’noi i Teoreticheskoi Fiziki* **87**, 348–365 (1984).
- [183] Giuseppe Carleo, Federico Becca, Marco Schiró, and Michele Fabrizio, Localization and glassy dynamics of many-body quantum systems, *Scientific reports* **2**, 243 (2012).
- [184] Tarun Grover and Matthew P. A. Fisher, Quantum disentangled liquids, *J. Stat. Mech.* **2014**, P10010 (2014).
- [185] M. Schiulaz and M. Müller, Ideal quantum glass transitions: Many-body localization without quenched disorder, in *AIP Conference Proceedings*, Vol. 1610 (AIP Publishing, 2014) pp. 11–23.
- [186] Wojciech De Roeck and François Huveneers, Can translation invariant systems exhibit a many-body localized phase? in *From Particle Systems to Partial Differential Equations II* (Springer, 2015) pp. 173–192.

- [187] S. S. Kondov, W. R. McGehee, W. Xu, and B. DeMarco, Disorder-induced localization in a strongly correlated atomic hubbard gas, *Phys. Rev. Lett.* **114**, 083002 (2015).
- [188] Mauro Schiulaz, Alessandro Silva, and Markus Müller, Dynamics in many-body localized quantum systems without disorder, *Phys. Rev. B* **91**, 184202 (2015).
- [189] James M Hickey, Sam Genway, and Juan P Garrahan, Signatures of many-body localisation in a system without disorder and the relation to a glass transition, *Journal of Statistical Mechanics: Theory and Experiment* **2016**, 054047 (2016).
- [190] N. Y. Yao, C. R. Laumann, J. I. Cirac, M. D. Lukin, and J. E. Moore, Quasi-many-body localization in translation-invariant systems, *Phys. Rev. Lett.* **117**, 240601 (2016).
- [191] L. M. Falicov and J. C. Kimball, Simple model for semiconductor-metal transitions: Smb_6 and transition-metal oxides, *Phys. Rev. Lett.* **22**, 997–999 (1969).
- [192] Christian Prosko, Shu-Ping Lee, and Joseph Maciejko, Simple \mathbb{Z}_2 lattice gauge theories at finite fermion density, *Phys. Rev. B* **96**, 205104 (2017).
- [193] A. Smith, J. Knolle, D. L. Kovrizhin, and R. Moessner, Disorder-free localization, *Phys. Rev. Lett.* **118**, 266601 (2017).
- [194] A. Smith, J. Knolle, R. Moessner, and D. L. Kovrizhin, Absence of ergodicity without quenched disorder: From quantum disentangled liquids to many-body localization, *Phys. Rev. Lett.* **119**, 176601 (2017).
- [195] Adam Smith, Johannes Knolle, Roderich Moessner, and Dmitry L. Kovrizhin, Dynamical localization in \mathbb{Z}_2 lattice gauge theories, *Phys. Rev. B* **97**, 245137 (2018).
- [196] Marlon Brenes, Marcello Dalmonte, Markus Heyl, and Antonello Scardicchio, Many-body localization dynamics from gauge invariance, *Phys. Rev. Lett.* **120**, 030601 (2018).
- [197] Adam Smith, Johannes Knolle, Roderich Moessner, and Dmitry L. Kovrizhin, Logarithmic spreading of out-of-time-ordered correlators without many-body localization, *Phys. Rev. Lett.* **123**, 086602 (2019).
- [198] Angelo Russomanno, Simone Notarnicola, Federica Maria Surace, Rosario Fazio, Marcello Dalmonte, and Markus Heyl, Homogeneous floquet time crystal protected by gauge invariance, *Phys. Rev. Research* **2**, 012003 (2020).
- [199] P. Karpov, R. Verdel, Y.-P. Huang, M. Schmitt, and M. Heyl, Disorder-free localization in an interacting 2d lattice gauge theory, *Phys. Rev. Lett.* **126**, 130401 (2021).

- [200] Wojciech De Roeck, Francois Huveneers, Markus Müller, and Mauro Schiulaz, Absence of many-body mobility edges, *Phys. Rev. B* **93**, 014203 (2016).
- [201] Luigi Amico, Rosario Fazio, Andreas Osterloh, and Vlatko Vedral, Entanglement in many-body systems, *Rev. Mod. Phys.* **80**, 517–576 (2008).
- [202] Nicolas Laflorencie, Quantum entanglement in condensed matter systems, *Phys. Rep.* **646**, 1–59 (2016).
- [203] Nicholas E. Sherman, Trithep Devakul, Matthew B. Hastings, and Rajiv R. P. Singh, Nonzero-temperature entanglement negativity of quantum spin models: Area law, linked cluster expansions, and sudden death, *Phys. Rev. E* **93**, 022128 (2016).
- [204] Karol Życzkowski, Paweł Horodecki, Anna Sanpera, and Maciej Lewenstein, Volume of the set of separable states, *Phys. Rev. A* **58**, 883–892 (1998).
- [205] G. Vidal and R. F. Werner, Computable measure of entanglement, *Phys. Rev. A* **65**, 032314 (2002).
- [206] Hassan Shapourian, Ken Shiozaki, and Shinsei Ryu, Partial time-reversal transformation and entanglement negativity in fermionic systems, *Phys. Rev. B* **95**, 165101 (2017).
- [207] Jens Eisert, Viktor Eisler, and Zoltán Zimborás, Entanglement negativity bounds for fermionic gaussian states, *Phys. Rev. B* **97**, 165123 (2018).
- [208] Asher Peres, Separability criterion for density matrices, *Phys. Rev. Lett.* **77**, 1413–1415 (1996).
- [209] J. Eisert, *Entanglement in quantum information theory*, Ph.D. thesis, University of Potsdam (2001), [arXiv:quant-ph/0610253v1](https://arxiv.org/abs/quant-ph/0610253v1) .
- [210] Hannu Wichterich, Julien Vidal, and Sougato Bose, Universality of the negativity in the Lipkin-Meshkov-Glick model, *Phys. Rev. A* **81**, 032311 (2010).
- [211] Raul A. Santos, V. Korepin, and Sougato Bose, Negativity for two blocks in the one-dimensional spin-1 Affleck-Kennedy-Lieb-Tasaki model, *Phys. Rev. A* **84**, 062307 (2011).
- [212] Raul A. Santos and V. Korepin, Negativity in the generalized valence bond solid state, *Quantum Inf. Process.* **15**, 4581–4598 (2016).
- [213] C. Castelnovo, Negativity and topological order in the toric code, *Phys. Rev. A* **88**, 042319 (2013).

- [214] Yirun Arthur Lee and Guifre Vidal, Entanglement negativity and topological order, *Phys. Rev. A* **88**, 042318 (2013).
- [215] Pasquale Calabrese, John Cardy, and Erik Tonni, Entanglement negativity in quantum field theory, *Phys. Rev. Lett.* **109**, 130502 (2012).
- [216] Pasquale Calabrese, John Cardy, and Erik Tonni, Entanglement negativity in extended systems: a field theoretical approach, *J. Stat. Mech.* **2013**, P02008 (2013).
- [217] M. C. Arnesen, S. Bose, and V. Vedral, Natural thermal and magnetic entanglement in the 1d Heisenberg model, *Phys. Rev. Lett.* **87**, 017901 (2001).
- [218] E. Mehran, S. Mahdavifar, and R. Jafari, Induced effects of the Dzyaloshinskii-Moriya interaction on the thermal entanglement in spin-1/2 Heisenberg chains, *Phys. Rev. A* **89**, 042306 (2014).
- [219] D. Gunlycke, V. M. Kendon, V. Vedral, and S. Bose, Thermal concurrence mixing in a one-dimensional Ising model, *Phys. Rev. A* **64**, 042302 (2001).
- [220] Tobias J. Osborne and Michael A. Nielsen, Entanglement in a simple quantum phase transition, *Phys. Rev. A* **66**, 032110 (2002).
- [221] Xiaoguang Wang, Hongchen Fu, and Allan I Solomon, Thermal entanglement in three-qubit Heisenberg models, *J. Phys. A: Math. Gen.* **34**, 11307 (2001).
- [222] Xiaoguang Wang and Paolo Zanardi, Quantum entanglement and Bell inequalities in Heisenberg spin chains, *Physics Letters A* **301**, 1–6 (2002).
- [223] Xiaoguang Wang, Thermal and ground-state entanglement in Heisenberg XX qubit rings, *Phys. Rev. A* **66**, 034302 (2002).
- [224] Guo-Feng Zhang and Shu-Shen Li, Thermal entanglement in a two-qubit Heisenberg XXZ spin chain under an inhomogeneous magnetic field, *Phys. Rev. A* **72**, 034302 (2005).
- [225] Da-Chuang Li, Xian-Ping Wang, and Zhuo-Liang Cao, Thermal entanglement in the anisotropic Heisenberg XXZ model with the Dzyaloshinskii-Moriya interaction, *J. Phys.: Condens. Matter* **20**, 325229 (2008).
- [226] L. Zhou, H. S. Song, Y. Q. Guo, and C. Li, Enhanced thermal entanglement in an anisotropic Heisenberg XYZ chain, *Phys. Rev. A* **68**, 024301 (2003).
- [227] Gustavo Rigolin, Thermal entanglement in the two-qubit Heisenberg XYZ model, *Int. J. Quantum Inform.* **02**, 393–405 (2004).

- [228] G. Lagmago Kamta and Anthony F. Starace, Anisotropy and magnetic field effects on the entanglement of a two qubit Heisenberg XY chain, *Phys. Rev. Lett.* **88**, 107901 (2002).
- [229] Xiaoguang Wang and Z. D. Wang, Thermal entanglement in ferrimagnetic chains, *Phys. Rev. A* **73**, 064302 (2006).
- [230] Ma Xiao-San and Wang An-Min, Thermal entanglement of an XY two-qutrit spin chain with dzialoshinski–moriya interaction, *Commun. Theor. Phys.* **52**, 825 (2009).
- [231] Yu-Liang Xu, Lu-Shun Wang, and Xiang-Mu Kong, Thermal entanglement between non-nearest-neighbor spins on fractal lattices, *Phys. Rev. A* **87**, 012312 (2013).
- [232] Viktor Eisler and Zoltán Zimborás, Entanglement negativity in the harmonic chain out of equilibrium, *New J. Phys.* **16**, 123020 (2014).
- [233] Pasquale Calabrese, John Cardy, and Erik Tonni, Finite temperature entanglement negativity in conformal field theory, *J. Phys. A: Math. Theor.* **48**, 015006 (2015).
- [234] Claudio Castelnovo and Claudio Chamon, Entanglement and topological entropy of the toric code at finite temperature, *Phys. Rev. B* **76**, 184442 (2007).
- [235] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki, Mixed-state entanglement and distillation: Is there a “bound” entanglement in nature? *Phys. Rev. Lett.* **80**, 5239–5242 (1998).
- [236] Pasquale Calabrese, Luca Tagliacozzo, and Erik Tonni, Entanglement negativity in the critical Ising chain, *J. Stat. Mech.* **2013**, P05002 (2013).
- [237] Paola Ruggiero, Vincenzo Alba, and Pasquale Calabrese, Negativity spectrum of one-dimensional conformal field theories, *Phys. Rev. B* **94**, 195121 (2016).
- [238] Claudio Castelnovo and Claudio Chamon, Quantum topological phase transition at the microscopic level, *Phys. Rev. B* **77**, 054433 (2008).
- [239] Claudio Castelnovo and Claudio Chamon, Topological order and topological entropy in classical systems, *Phys. Rev. B* **76**, 174416 (2007).
- [240] Alioscia Hamma, Radu Ionicioiu, and Paolo Zanardi, Bipartite entanglement and entropic boundary law in lattice spin systems, *Phys. Rev. A* **71**, 022315 (2005).
- [241] O. Hart and C. Castelnovo, Entanglement negativity and sudden death in the toric code at finite temperature, *Phys. Rev. B* **97**, 144410 (2018).

- [242] Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki, Separability of mixed states: necessary and sufficient conditions, *Phys. Lett. A* **223** (1996).
- [243] C. N. Yang and T. D. Lee, Statistical theory of equations of state and phase transitions. i. theory of condensation, *Phys. Rev.* **87**, 404–409 (1952).
- [244] T. D. Lee and C. N. Yang, Statistical theory of equations of state and phase transitions. ii. lattice gas and ising model, *Phys. Rev.* **87**, 410–419 (1952).
- [245] Alessandro Ferraro, Daniel Cavalcanti, Artur García-Saez, and Antonio Acín, Thermal bound entanglement in macroscopic systems and area law, *Phys. Rev. Lett.* **100**, 080502 (2008).
- [246] Leon Balents, Spin liquids in frustrated magnets, *Nature* **464**, 199–208 (2010).
- [247] Tian-Heng Han, Joel S Helton, Shaoyan Chu, Daniel G Nocera, Jose A Rodriguez-Rivera, Collin Broholm, and Young S Lee, Fractionalized excitations in the spin-liquid state of a kagome-lattice antiferromagnet, *Nature* **492**, 406 (2012).
- [248] Yao Shen *et al.*, Evidence for a spinon fermi surface in a triangular-lattice quantum-spin-liquid candidate, *Nature* **540**, 559–562.
- [249] Joseph A. M. Paddison, Marcus Daum, Zhiling Dun, Georg Ehlers, Yaohua Liu, Matthew B. Stone, Haidong Zhou, and Martin Mourigal, Continuous excitations of the triangular-lattice quantum spin liquid YbMgGaO₄, *Nature Physics* **13**, 117–122.
- [250] X. G. Wen and Q. Niu, Ground-state degeneracy of the fractional quantum hall states in the presence of a random potential and on high-genus riemann surfaces, *Phys. Rev. B* **41**, 9377–9396 (1990).
- [251] Joji Nasu, Johannes Knolle, Dima L Kovrizhin, Yukitoshi Motome, and Roderich Moessner, Fermionic response from fractionalization in an insulating two-dimensional magnet, *Nat. Phys.* **12**, 912 (2016).
- [252] Junki Yoshitake, Joji Nasu, and Yukitoshi Motome, Fractional Spin Fluctuations as a Precursor of Quantum Spin Liquids: Majorana Dynamical Mean-Field Study for the Kitaev Model, *Phys. Rev. Lett.* **117**, 157203 (2016).
- [253] Junki Yoshitake, Joji Nasu, Yasuyuki Kato, and Yukitoshi Motome, Majorana dynamical mean-field study of spin dynamics at finite temperatures in the honeycomb Kitaev model, *Phys. Rev. B* **96**, 024438 (2017).

- [254] Junki Yoshitake, Joji Nasu, and Yukiotoshi Motome, Temperature evolution of spin dynamics in two- and three-dimensional Kitaev models: Influence of fluctuating Z_2 flux, *Phys. Rev. B* **96**, 064433 (2017).
- [255] Gábor B. Halász, Stefanos Kourtis, Johannes Knolle, and Natalia B. Perkins, Observing spin fractionalization in the Kitaev spin liquid via temperature evolution of indirect resonant inelastic x-ray scattering, *Phys. Rev. B* **99**, 184417 (2019).
- [256] Roderich Moessner and Kumar S. Raman, Quantum dimer models, in *Introduction to Frustrated Magnetism: Materials, Experiments, Theory*, edited by Claudine Lacroix, Philippe Mendels, and Frédéric Mila (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011) pp. 437–479.
- [257] J. M. Luttinger and W. Kohn, Motion of electrons and holes in perturbed periodic fields, *Phys. Rev.* **97**, 869–883 (1955).
- [258] J. R. Schrieffer and P. A. Wolff, Relation between the anderson and kondo hamiltonians, *Phys. Rev.* **149**, 491–492 (1966).
- [259] Jürgen Stein, Flow equations and the strong-coupling expansion for the hubbard model, *Journal of Statistical Physics* **88**, 487–511 (1997).
- [260] Christian Knetter and Goetz S Uhrig, Perturbation theory by flow equations: dimerized and frustrated $S = 1/2$ chain, *The European Physical Journal B-Condensed Matter and Complex Systems* **13**, 209–225 (2000).
- [261] Christian Knetter, Kai P Schmidt, and G tz S Uhrig, The structure of operators in effective particle-conserving models, *Journal of Physics A: Mathematical and General* **36**, 7889–7907 (2003).
- [262] Renate Gade, Anderson localization for sublattice models, *Nuclear Physics B* **398**, 499 – 515 (1993).
- [263] Andreas W. W. Ludwig, Matthew P. A. Fisher, R. Shankar, and G. Grinstein, Integer quantum hall transition: An alternative approach and exact results, *Phys. Rev. B* **50**, 7526–7552 (1994).
- [264] Olexei Motrunich, Kedar Damle, and David A. Huse, Particle-hole symmetric localization in two dimensions, *Phys. Rev. B* **65**, 064206 (2002).
- [265] W. F. Brinkman and T. M. Rice, Single-particle excitations in magnetic insulators, *Phys. Rev. B* **2**, 1324–1338 (1970).

- [266] M M Mohan, Possibility of quasiparticle behaviour in the strongly correlated hubbard model, *Journal of Physics: Condensed Matter* **3**, 4307–4312 (1991).
- [267] Oleg A. Starykh and George F. Reiter, Hole motion in the ising antiferromagnet: An application of the recursion method, *Phys. Rev. B* **53**, 2517–2522 (1996).
- [268] A. L. Chernyshev and P. W. Leung, Holes in the $t - J_z$ model: A diagrammatic study, *Phys. Rev. B* **60**, 1592–1606 (1999).
- [269] B. L. Altshuler and L. B. Ioffe, Motion of fast particles in strongly fluctuating magnetic fields, *Phys. Rev. Lett.* **69**, 2979–2982 (1992).
- [270] John B. Kogut, An introduction to lattice gauge theory and spin systems, *Rev. Mod. Phys.* **51**, 659–713 (1979).
- [271] Hans De Raedt, Product formula algorithms for solving the time dependent schrödinger equation, *Computer Physics Reports* **7**, 1–72 (1987).
- [272] Olga Petrova, Roderich Moessner, and S. L. Sondhi, Hydrogenic states of monopoles in diluted quantum spin ice, *Phys. Rev. B* **92**, 100401 (2015).
- [273] Masafumi Udagawa and Roderich Moessner, Spectrum of itinerant fractional excitations in quantum spin ice, *Phys. Rev. Lett.* **122**, 117201 (2019).
- [274] S. A. Trugman, Interaction of holes in a hubbard antiferromagnet and high-temperature superconductivity, *Phys. Rev. B* **37**, 1597–1603 (1988).
- [275] Ian. M. Wanless, Counting matchings and tree-like walks in regular graphs, *Combinatorics, Probability and Computing* **19**, 463–480 (2010).
- [276] H. N. V. Temperley, Combinatorial problems suggested by the statistical mechanics of domains and of rubber-like molecules, *Phys. Rev.* **103**, 1–16 (1956).
- [277] Tohru Morita, Useful procedure for computing the lattice green’s function-square, tetragonal, and bcc lattices, *Journal of Mathematical Physics* **12**, 1744–1747 (1971).
- [278] Alexander Altland and B D Simons, Field theory of the random flux model, *Journal of Physics A: Mathematical and General* **32**, L353 (1999).
- [279] Alexander Altland and B.D. Simons, Field theory of the random flux model, *Nuclear Physics B* **562**, 445 – 476 (1999).
- [280] Akira Furusaki, Anderson localization due to a random magnetic field in two dimensions, *Phys. Rev. Lett.* **82**, 604–607 (1999).

- [281] Joji Nasu, Masafumi Udagawa, and Yukiotoshi Motome, Thermal fractionalization of quantum spins in a kitaev model: Temperature-linear specific heat and coherent transport of majorana fermions, *Phys. Rev. B* **92**, 115122 (2015).
- [282] Márton Kanász-Nagy, Izabella Lovas, Fabian Grusdt, Daniel Greif, Markus Greiner, and Eugene A. Demler, Quantum correlations at infinite temperature: The dynamical nagaoka effect, *Phys. Rev. B* **96**, 014303 (2017).
- [283] Tohru Kawarabayashi and Tomi Ohtsuki, Diffusion of electrons in random magnetic fields, *Phys. Rev. B* **51**, 10897–10904 (1995).
- [284] Athmane Tadjine and Christophe Delerue, Anderson localization induced by gauge-invariant bond-sign disorder in square PbSe nanocrystal lattices, *Phys. Rev. B* **98**, 125412 (2018).
- [285] H. Takagi, T. Takayama, G. Jackeli, G. Khaliullin, and S. E. Nagler, Kitaev quantum spin liquid - concept and materialization, [arXiv:1903.08081](https://arxiv.org/abs/1903.08081) (2019).
- [286] Claudio Chamon, Dmitry Green, and Zhi-Cheng Yang, How to emulate quantum spin liquids and build topological qubits with available quantum hardware, (2019), [arXiv:1908.04791](https://arxiv.org/abs/1908.04791) .
- [287] Johan Carlström, Nikolay Prokof'ev, and Boris Svistunov, Quantum walk in degenerate spin environments, *Phys. Rev. Lett.* **116**, 247202 (2016).
- [288] A Bohrdt, F Grusdt, and M Knap, Dynamical formation of a magnetic polaron in a two-dimensional quantum antiferromagnet, *New Journal of Physics* **22**, 123023 (2020).
- [289] N V Prokof'ev and P C E Stamp, Theory of the spin bath, *Reports on Progress in Physics* **63**, 669–726 (2000).
- [290] N. V. Prokof'ev and P. C. E. Stamp, Decoherence and quantum walks: Anomalous diffusion and ballistic tails, *Phys. Rev. A* **74**, 020102 (2006).
- [291] Waseem S. Bakr, Jonathon I. Gillen, Amy Peng, Simon Fölling, and Markus Greiner, A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice, *Nature* **462**, 74–77 (2009).
- [292] Jacob F. Sherson, Christof Weitenberg, Manuel Endres, Marc Cheneau, Immanuel Bloch, and Stefan Kuhr, Single-atom-resolved fluorescence imaging of an atomic mott insulator, *Nature* **467**, 68–72 (2010).
- [293] Elmar Haller, James Hudson, Andrew Kelly, Dylan A. Cotta, Bruno Peaudecerf, Graham D. Bruce, and Stefan Kuhr, Single-atom imaging of fermions in a quantum-gas microscope, *Nature Physics* **11**, 738–742 (2015).

- [294] Lawrence W. Cheuk, Matthew A. Nichols, Melih Okan, Thomas Gersdorf, Vinay V. Ramasesh, Waseem S. Bakr, Thomas Lompe, and Martin W. Zwierlein, Quantum-gas microscope for fermionic atoms, *Phys. Rev. Lett.* **114**, 193001 (2015).
- [295] Maxwell F. Parsons, Florian Huber, Anton Mazurenko, Christie S. Chiu, Widagdo Setiawan, Katherine Wooley-Brown, Sebastian Blatt, and Markus Greiner, Site-resolved imaging of fermionic ${}^6\text{Li}$ in an optical lattice, *Phys. Rev. Lett.* **114**, 213002 (2015).
- [296] G. J. A. Edge, R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky, and J. H. Thywissen, Imaging and addressing of individual fermionic atoms in an optical lattice, *Phys. Rev. A* **92**, 063406 (2015).
- [297] Michael A. Levin and Xiao-Gang Wen, String-net condensation: A physical mechanism for topological phases, *Phys. Rev. B* **71**, 045110 (2005).
- [298] Oliver Hart, Yuan Wan, and Claudio Castelnovo, Correlation holes and slow dynamics induced by fractional statistics in gapped quantum spin liquids, *Nature communications* **12**, 1–8 (2021).
- [299] Yosuke Nagaoka, Ground state of correlated electrons in a narrow almost half-filled s band, *Solid State Communications* **3**, 409–412 (1965).
- [300] Yosuke Nagaoka, Ferromagnetism in a narrow, almost half-filled s band, *Phys. Rev.* **147**, 392–405 (1966).
- [301] Steven T. Bramwell and Michel J. P. Gingras, Spin ice state in frustrated magnetic pyrochlore materials, *Science* **294**, 1495–1501 (2001).
- [302] C. Castelnovo, R. Moessner, and S. L. Sondhi, Thermal quenches in spin ice, *Phys. Rev. Lett.* **104**, 107201 (2010).
- [303] Carley Paulsen, SR Giblin, Elsa Lhotel, D Prabhakaran, Geetha Balakrishnan, K Matsuhira, and ST Bramwell, Experimental signature of the attractive coulomb force between positive and negative magnetic monopoles in spin ice, *Nature Physics* **12**, 661 (2016).
- [304] Carley Paulsen, Martin James Jackson, Elsa Lhotel, Benjamin Canals, D Prabhakaran, K Matsuhira, SR Giblin, and ST Bramwell, Far-from-equilibrium monopole dynamics in spin ice, *Nature Physics* **10**, 135 (2014).
- [305] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, 1999).

- [306] David P. Landau and Kurt Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*, 4th ed. (Cambridge University Press, 2014).
- [307] Werner Krauth, *Statistical mechanics: algorithms and computations*, Vol. 13 (OUP Oxford, 2006).
- [308] Hidemaro Suwa and Syngae Todo, Markov chain monte carlo method without detailed balance, *Phys. Rev. Lett.* **105**, 120603 (2010).
- [309] Sebastian C. Kapfer and Werner Krauth, Irreversible local markov chains with rapid convergence towards equilibrium, *Phys. Rev. Lett.* **119**, 240603 (2017).
- [310] Ulli Wolff, Critical slowing down, *Nuclear Physics B - Proceedings Supplements* **17**, 93–102 (1990).
- [311] Ulli Wolff, Collective monte carlo updating for spin systems, *Phys. Rev. Lett.* **62**, 361–364 (1989).
- [312] J. Snyder, B. G. Ueland, J. S. Slusky, H. Karunadasa, R. J. Cava, and P. Schiffer, Low-temperature spin freezing in the $\text{Dy}_2\text{Ti}_2\text{O}_7$ spin ice, *Phys. Rev. B* **69**, 064414 (2004).
- [313] Jesper Dall and Paolo Sibani, Faster monte carlo simulations at low temperatures. the waiting time method, *Computer Physics Communications* **141**, 260 – 267 (2001).
- [314] Jesper Dall and Paolo Sibani, Exploring valleys of aging systems: the spin glass case, *The European Physical Journal B-Condensed Matter and Complex Systems* **36**, 233–243 (2003).
- [315] A.B. Bortz, M.H. Kalos, and J.L. Lebowitz, A new algorithm for monte carlo simulation of ising spin systems, *Journal of Computational Physics* **17**, 10–18 (1975).
- [316] Sarah Mostame, Claudio Castellano, Roderich Moessner, and Shivaji L. Sondhi, Tunable nonequilibrium dynamics of field quenches in spin ice, *Proceedings of the National Academy of Sciences* **111**, 640–645 (2014).
- [317] A. A. Ovchinnikov and V. V. Atrazhev, The role of space fluctuations in kinetics of recombination of charged particles, *Physica A: Statistical Mechanics and its Applications* **276**, 1 – 9 (2000).
- [318] Simon W de Leeuw, John William Perram, and Edgar Roderick Smith, Simulation of electrostatic systems in periodic boundary conditions. i. lattice sums and dielectric constants, *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **373**, 27–56 (1980).

- [319] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*, Computational science (Elsevier Science, 2001).
- [320] Christopher L. Henley, The “coulomb phase” in frustrated systems, *Annual Review of Condensed Matter Physics* **1**, 179–210 (2010).
- [321] Valeriy V. Ginzburg, Leo Radzihovsky, and Noel A. Clark, Self-consistent model of an annihilation-diffusion reaction with long-range interactions, *Phys. Rev. E* **55**, 395–402 (1997).
- [322] Taras Yavors’kii, Tom Fennell, Michel J. P. Gingras, and Steven T. Bramwell, Dy₂Ti₂O₇ Spin Ice: A Test Case for Emergent Clusters in a Frustrated Magnet, *Phys. Rev. Lett.* **101**, 037204 (2008).
- [323] P. Henelius, T. Lin, M. Enjalran, Z. Hao, J. G. Rau, J. Altsaar, F. Flicker, T. Yavors’kii, and M. J. P. Gingras, Refrustration and competing orders in the prototypical Dy₂Ti₂O₇ spin ice material, *Phys. Rev. B* **93**, 024402 (2016).
- [324] R. A. Borzi, F. A. Gómez Albarracín, H. D. Rosales, G. L. Rossini, A. Steppke, D. Prabhakaran, A. P. Mackenzie, D. C. Cabra, and S. A. Grigera, Intermediate magnetization state and competing orders in Dy₂Ti₂O₇ and Ho₂Ti₂O₇, *Nature Communications* **7**, 12592 (2016).
- [325] Anjana M Samarakoon, Kipton Barros, Ying Wai Li, Markus Eisenbach, Qiang Zhang, Feng Ye, Z. L. Dun, Haidong Zhou, Santiago A. Grigera, Cristian D. Batista, and D. Alan Tennant, Machine Learning Assisted Insight to Spin Ice Dy₂Ti₂O₇, (2019), [arXiv:1906.11275](https://arxiv.org/abs/1906.11275).
- [326] Attila Szabó and Claudio Castelnovo, Seeing beyond the light: Vison and photon electrodyamics in quantum spin ice, *Phys. Rev. B* **100**, 014417 (2019).
- [327] D. Slobinsky, C. Castelnovo, R. A. Borzi, A. S. Gibbs, A. P. Mackenzie, R. Moessner, and S. A. Grigera, Unconventional Magnetization Processes and Thermal Runaway in Spin-Ice Dy₂Ti₂O₇, *Phys. Rev. Lett.* **105**, 267205 (2010).
- [328] Clemens Müller, Jared H Cole, and Jürgen Lisenfeld, Towards understanding two-level-systems in amorphous solids: insights from quantum circuits, *Reports on Progress in Physics* **82**, 124501 (2019).
- [329] Juan Maldacena, Stephen H. Shenker, and Douglas Stanford, A bound on chaos, *Journal of High Energy Physics* **2016**, 106 (2016).
- [330] D.M. Basko, I.L. Aleiner, and B.L. Altshuler, Metal–insulator transition in a weakly interacting many-electron system with localized single-particle states, *Ann. Phys. (N.Y.)* **321**, 1126–1205 (2006).

- [331] Sarang Gopalakrishnan and S.A. Parameswaran, Dynamics and transport at the threshold of many-body localization, *Physics Reports* **862**, 1 – 62 (2020).
- [332] Yu Kagan and LA Maksimov, Localization in a system of interacting particles diffusing in a regular crystal, *Zh. Eksp. Teor. Fiz.* **87**, 348–365 (1984).
- [333] Wojciech De Roeck and François Huveneers, Asymptotic quantum many-body localization from thermal disorder, *Communications in Mathematical Physics* **332**, 1017–1082 (2014).
- [334] Merlijn van Horssen, Emanuele Levi, and Juan P. Garrahan, Dynamics of many-body localization in a translation-invariant quantum glass model, *Phys. Rev. B* **92**, 100305 (2015).
- [335] Z. Papić, E. Miles Stoudenmire, and Dmitry A. Abanin, Many-body localization in disorder-free systems: The importance of finite-size constraints, *Annals of Physics* **362**, 714 – 725 (2015).
- [336] Sarang Gopalakrishnan and David A. Huse, Instability of many-body localized systems as a phase transition in a nonstandard thermodynamic limit, *Phys. Rev. B* **99**, 134305 (2019).
- [337] S. A. Parameswaran and S. Gopalakrishnan, Non-fermi glasses: Localized descendants of fractionalized metals, *Phys. Rev. Lett.* **119**, 146601 (2017).
- [338] Michael Schreiber, Sean S. Hodgman, Pranjal Bordia, Henrik P. Lüschen, Mark H. Fischer, Ronen Vosk, Ehud Altman, Ulrich Schneider, and Immanuel Bloch, Observation of many-body localization of interacting fermions in a quasirandom optical lattice, *Science* **349**, 842–845 (2015).
- [339] Jae-yoon Choi, Sebastian Hild, Johannes Zeiher, Peter Schauß, Antonio Rubio-Abadal, Tarik Yefsah, Vedika Khemani, David A. Huse, Immanuel Bloch, and Christian Gross, Exploring the many-body localization transition in two dimensions, *Science* **352**, 1547–1552 (2016).
- [340] Pranjal Bordia, Henrik P. Lüschen, Sean S. Hodgman, Michael Schreiber, Immanuel Bloch, and Ulrich Schneider, Coupling identical one-dimensional many-body localized systems, *Phys. Rev. Lett.* **116**, 140401 (2016).
- [341] Jacob Smith, Aaron Lee, Philip Richerme, Brian Neyenhuis, Paul W Hess, Philipp Hauke, Markus Heyl, David A Huse, and Christopher Monroe, Many-body localization in a quantum simulator with programmable random disorder, *Nature Physics* **12**, 907 (2016).
- [342] Alexander Lukin, Matthew Rispoli, Robert Schittko, M. Eric Tai, Adam M. Kaufman, Soonwon Choi, Vedika Khemani, Julian Léonard, and Markus Greiner, Probing entanglement in a many-body-localized system, *Science* **364**, 256–260 (2019).

- [343] Ben Chiaro *et al.*, Growth and preservation of entanglement in a many-body localized system, arXiv preprint arXiv:1910.06024 (2019).
- [344] M. Ovadia, B. Sacépé, and D. Shahar, Electron-phonon decoupling in disordered insulators, *Phys. Rev. Lett.* **102**, 176802 (2009).
- [345] Z. Ovadyahu, Suppression of inelastic electron-electron scattering in anderson insulators, *Phys. Rev. Lett.* **108**, 156602 (2012).
- [346] Z. Ovadyahu, Thermalization processes in interacting anderson insulators, *Phys. Rev. B* **91**, 035113 (2015).
- [347] M Ovadia, D Kalok, I Tamir, S Mitra, B Sacépé, and D Shahar, Evidence for a finite-temperature insulator, *Scientific reports* **5**, 13503 (2015).
- [348] Thanh Nguyen, Nina Andrejevic, Hoi Chun Po, Yoichiro Tsurimaki, Nathan C Drucker, Ahmet Alatas, Ercan E Alp, Bogdan M Leu, Alessandro Cunsolo, Yong Q Cai, *et al.*, Signature of many-body localization of phonons in strongly disordered superlattices, arXiv preprint arXiv:2008.02257 (2020).
- [349] Wojciech Brzezicki and Andrzej M. Oleś, Exact solution for a quantum compass ladder, *Phys. Rev. B* **80**, 014405 (2009).
- [350] M. Serbyn, M. Knap, S. Gopalakrishnan, Z. Papić, N. Y. Yao, C. R. Laumann, D. A. Abanin, M. D. Lukin, and E. A. Demler, Interferometric probes of many-body localization, arXiv:1403.0693 (2014).
- [351] Maksym Serbyn, Z. Papić, and D. A. Abanin, Quantum quenches in the many-body localized phase, *Phys. Rev. B* **90**, 174302 (2014).
- [352] Dong-Ling Deng, Xiaopeng Li, J. H. Pixley, Yang-Le Wu, and S. Das Sarma, Logarithmic entanglement lightcone in many-body localized systems, *Phys. Rev. B* **95**, 024202 (2017).
- [353] R. Vasseur, S. A. Parameswaran, and J. E. Moore, Quantum revivals and many-body localization, *Phys. Rev. B* **91**, 140202 (2015).
- [354] Sarang Gopalakrishnan, Markus Müller, Vedika Khemani, Michael Knap, Eugene Demler, and David A. Huse, Low-frequency conductivity in many-body localized systems, *Phys. Rev. B* **92**, 104202 (2015).
- [355] Elbio Dagotto, Experiments on ladders reveal a complex interplay between a spin-gapped normal state and superconductivity, *Reports on Progress in Physics* **62**, 1525–1571 (1999).

- [356] Loredana M. Vasiloiu, Federico Carollo, Matteo Marcuzzi, and Juan P. Garrahan, Strong zero modes in a class of generalized ising spin ladders with plaquette interactions, *Phys. Rev. B* **100**, 024309 (2019).
- [357] Desmond A Johnston and Ranasinghe PKCM Ranasinghe, (four) dual plaquette 3d ising models, *Entropy* **22**, 633 (2020).
- [358] Rahul M. Nandkishore and Michael Hermele, Fractons, *Annual Review of Condensed Matter Physics* **10**, 295–313 (2019).
- [359] Abhinav Prem, Jeongwan Haah, and Rahul Nandkishore, Glassy quantum dynamics in translation invariant fracton models, *Phys. Rev. B* **95**, 155133 (2017).
- [360] Cenke Xu and J. E. Moore, Strong-weak coupling self-duality in the two-dimensional quantum phase transition of $p + ip$ superconducting arrays, *Phys. Rev. Lett.* **93**, 047003 (2004).
- [361] Cenke Xu and J.E. Moore, Reduction of effective dimensionality in lattice models of superconducting arrays and frustrated magnets, *Nuclear Physics B* **716**, 487–508 (2005).
- [362] Zohar Nussinov and Eduardo Fradkin, Discrete sliding symmetries, dualities, and self-dualities of quantum orbital compass models and $p + ip$ superconducting arrays, *Phys. Rev. B* **71**, 195120 (2005).
- [363] Roger A. Horn and Charles R. Johnson, *Matrix Analysis* (Cambridge University Press, 1985).
- [364] David A. Huse, Rahul Nandkishore, Vadim Oganesyan, Arijeet Pal, and S. L. Sondhi, Localization-protected quantum order, *Phys. Rev. B* **88**, 014206 (2013).
- [365] David Pekker, Gil Refael, Ehud Altman, Eugene Demler, and Vadim Oganesyan, Hilbert-glass transition: New universality of temperature-tuned many-body dynamical quantum criticality, *Phys. Rev. X* **4**, 011052 (2014).
- [366] Jonas A. Kjäll, Jens H. Bardarson, and Frank Pollmann, Many-body localization in a disordered quantum ising chain, *Phys. Rev. Lett.* **113**, 107204 (2014).
- [367] Daniel S. Fisher, Critical behavior of random transverse-field ising spin chains, *Phys. Rev. B* **51**, 6411–6461 (1995).
- [368] Maurizio Fagotti and Pasquale Calabrese, Entanglement entropy of two disjoint blocks in XY chains, *Journal of Statistical Mechanics: Theory and Experiment* **2010**, P04016 (2010).

- [369] Elliott Lieb, Theodore Schultz, and Daniel Mattis, Two soluble models of an antiferromagnetic chain, *Annals of Physics* **16**, 407 – 466 (1961).
- [370] J.H.H. Perk, H.W. Capel, G.R.W. Quispel, and F.W. Nijhoff, Finite-temperature correlations for the ising chain in a transverse field, *Physica A: Statistical Mechanics and its Applications* **123**, 1 – 49 (1984).
- [371] Shreya Vardhan, Giuseppe De Tomasi, Markus Heyl, Eric J. Heller, and Frank Pollmann, Characterizing time irreversibility in disordered fermionic systems by the effect of local perturbations, *Phys. Rev. Lett.* **119**, 016802 (2017).
- [372] Max McGinley, Andreas Nunnenkamp, and Johannes Knolle, Slow growth of out-of-time-order correlators and entanglement entropy in integrable disordered systems, *Phys. Rev. Lett.* **122**, 020603 (2019).
- [373] Giuseppe De Tomasi, Sthitadhi Roy, and Soumya Bera, Generalized dyson model: Nature of the zero mode and its implication in dynamics, *Phys. Rev. B* **94**, 144202 (2016).
- [374] Bernd Braunecker, Response of a fermi gas to time-dependent perturbations: Riemann-hilbert approach at nonzero temperatures, *Phys. Rev. B* **73**, 075122 (2006).
- [375] Bruno Tomasello, Claudio Castelnovo, Roderich Moessner, and Jorge Quintanilla, Correlated quantum tunneling of monopoles in spin ice, *Phys. Rev. Lett.* **123**, 067204 (2019).
- [376] Matthew Stern, Claudio Castelnovo, Roderich Moessner, Vadim Oganesyan, and Sarang Gopalakrishnan, Quantum percolation of monopole paths and the response of quantum spin ice, (2019), [arXiv:1911.05742 \[cond-mat.stat-mech\]](https://arxiv.org/abs/1911.05742) .
- [377] Yoshifumi Tokiwa, Takuya Yamashita, Daiki Terazawa, Kenta Kimura, Yuichi Kasahara, Takafumi Onishi, Yasuyuki Kato, Mario Halim, Philipp Gegenwart, Takasada Shibauchi, Satoru Nakatsuji, Eun-Gook Moon, and Yuji Matsuda, Discovery of emergent photon and monopoles in a quantum spin liquid, *Journal of the Physical Society of Japan* **87**, 064702 (2018).
- [378] F. Schmitt, P. S. Kirchmann, U. Bovensiepen, R. G. Moore, L. Rettig, M. Krenz, J.-H. Chu, N. Ru, L. Perfetti, D. H. Lu, M. Wolf, I. R. Fisher, and Z.-X. Shen, Transient Electronic Structure and Melting of a Charge Density Wave in TbTe_3 , *Science* **321**, 1649–1652 (2008).
- [379] W Kuehn, K Reimann, M Woerner, T Elsaesser, and R Hey, Two-dimensional terahertz correlation spectra of electronic excitations in semiconductor quantum wells, *The Journal of Physical Chemistry B* **115**, 5448–5455 (2011).

- [380] Michael Woerner, Wilhelm Kuehn, Pamela Bowlan, Klaus Reimann, and Thomas Elsaesser, Ultrafast two-dimensional terahertz spectroscopy of elementary excitations in solids, *New Journal of Physics* **15**, 025039 (2013).
- [381] Margherita Maiuri, Marco Garavelli, and Giulio Cerullo, Ultrafast spectroscopy: State of the art and open challenges, *Journal of the American Chemical Society* **142**, 3–15 (2019).
- [382] D. Jaksch and P. Zoller, The cold atom hubbard toolbox, *Annals of Physics* **315**, 52–79 (2005), special Issue.
- [383] Maciej Lewenstein, Anna Sanpera, Veronica Ahufinger, Bogdan Damski, Aditi Sen(De), and Ujjwal Sen, Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond, *Advances in Physics* **56**, 243–379 (2007).
- [384] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger, Many-body physics with ultracold gases, *Rev. Mod. Phys.* **80**, 885–964 (2008).
- [385] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbene, Quantum simulations with ultracold quantum gases, *Nature Physics* **8**, 267–276 (2012).
- [386] Christian Gross and Immanuel Bloch, Quantum simulations with ultracold atoms in optical lattices, *Science* **357**, 995–1001 (2017).
- [387] Rainer Blatt and Christian F Roos, Quantum simulations with trapped ions, *Nature Physics* **8**, 277–284 (2012).
- [388] Ch Schneider, Diego Porras, and Tobias Schaetz, Experimental quantum simulations of many-body physics with trapped ions, *Reports on Progress in Physics* **75**, 024401 (2012).
- [389] Jiehang Zhang, Guido Pagano, Paul W Hess, Antonis Kyprianidis, Patrick Becker, Harvey Kaplan, Alexey V Gorshkov, Z-X Gong, and Christopher Monroe, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, *Nature* **551**, 601–604 (2017).
- [390] Hendrik Weimer, Markus Müller, Igor Lesanovsky, Peter Zoller, and Hans Peter Büchler, A rydberg quantum simulator, *Nature Physics* **6**, 382–388 (2010).
- [391] Andrew A Houck, Hakan E Türeci, and Jens Koch, On-chip quantum simulation with superconducting circuits, *Nature Physics* **8**, 292–299 (2012).
- [392] Gheorghe-Sorin Paraoanu, Recent progress in quantum simulation using superconducting circuits, *Journal of Low Temperature Physics* **175**, 633–654 (2014).

- [393] Iulia Buluta and Franco Nori, Quantum simulators, *Science* **326**, 108–111 (2009).
- [394] I. M. Georgescu, S. Ashhab, and Franco Nori, Quantum simulation, *Rev. Mod. Phys.* **86**, 153–185 (2014).
- [395] K. J. Satzinger *et al.*, Realizing topologically ordered states on a quantum processor, (2021), [arXiv:2104.01180 \[quant-ph\]](https://arxiv.org/abs/2104.01180) .
- [396] Giulia Semeghini *et al.*, Probing topological spin liquids on a programmable quantum simulator, (2021), [arXiv:2104.04119 \[quant-ph\]](https://arxiv.org/abs/2104.04119) .
- [397] Chao-Yang Lu, Wei-Bo Gao, Otfried Gühne, Xiao-Qi Zhou, Zeng-Bing Chen, and Jian-Wei Pan, Demonstrating anyonic fractional statistics with a six-qubit quantum simulator, *Phys. Rev. Lett.* **102**, 030502 (2009).
- [398] J K Pachos, W Wieczorek, C Schmid, N Kiesel, R Pohlner, and H Weinfurter, Revealing anyonic features in a toric code quantum simulation, *New J. Phys.* **11**, 083010 (2009).
- [399] Annie Ji Hyun Park, Emma McKay, Dawei Lu, and Raymond Laflamme, Simulation of anyonic statistics and its topological path independence using a seven-qubit quantum simulator, *New Journal of Physics* **18**, 043043 (2016).
- [400] Y. P. Zhong, D. Xu, P. Wang, C. Song, Q. J. Guo, W. X. Liu, K. Xu, B. X. Xia, C.-Y. Lu, Siyuan Han, Jian-Wei Pan, and H. Wang, Emulating anyonic fractional statistical behavior in a superconducting quantum circuit, *Phys. Rev. Lett.* **117**, 110501 (2016).
- [401] Han-Ning Dai, Bing Yang, Andreas Reingruber, Hui Sun, Xiao-Fan Xu, Yu-Ao Chen, Zhen-Sheng Yuan, and Jian-Wei Pan, Four-body ring-exchange interactions and anyonic statistics within a minimal toric-code hamiltonian, *Nature Physics* **13**, 1195–1200 (2017).
- [402] Zhihuang Luo, Jun Li, Zhaokai Li, Ling-Yan Hung, Yidun Wan, Xinhua Peng, and Jiangfeng Du, Experimentally probing topological order and its breakdown through modular matrices, *Nature Physics* **14**, 160–165 (2018).
- [403] Chao Song, Da Xu, Pengfei Zhang, Jianwen Wang, Qiujiang Guo, Wuxin Liu, Kai Xu, Hui Deng, Keqiang Huang, Dongning Zheng, Shi-Biao Zheng, H. Wang, Xiaobo Zhu, Chao-Yang Lu, and Jian-Wei Pan, Demonstration of topological robustness of anyonic braiding statistics with a superconducting quantum circuit, *Phys. Rev. Lett.* **121**, 030502 (2018).
- [404] Christian Kraglund Andersen, Ants Remm, Stefania Lazar, Sebastian Krinner, Nathan Lacroix, Graham J Norris, Mihai Gabureac, Christopher Eichler, and Andreas Wallraff, Repeated quantum error detection in a surface code, *Nature Physics* **16**, 875–880 (2020).

- [405] Alexander Erhard, Hendrik Poulsen Nautrup, Michael Meth, Lukas Postler, Roman Stricker, Martin Stadler, Vlad Negnevitsky, Martin Ringbauer, Philipp Schindler, Hans J Briegel, *et al.*, Entangling logical qubits with lattice surgery, *Nature* **589**, 220–224 (2021).
- [406] Jian Lu, Xian Li, Harold Y. Hwang, Benjamin K. Ofori-Okai, Takayuki Kurihara, Tohru Suemoto, and Keith A. Nelson, Coherent two-dimensional terahertz magnetic resonance spectroscopy of collective spin waves, *Phys. Rev. Lett.* **118**, 207204 (2017).
- [407] Yuan Wan and N. P. Armitage, Resolving continua of fractional excitations by spinon echo in thz 2d coherent spectroscopy, *Phys. Rev. Lett.* **122**, 257401 (2019).
- [408] Fahad Mahmood, Dipanjan Chaudhuri, Sarang Gopalakrishnan, Rahul Nandkishore, and NP Armitage, Observation of a marginal fermi glass, *Nature Physics* , 1–5 (2021).
- [409] Wonjune Choi, Ki Hoon Lee, and Yong Baek Kim, Theory of two-dimensional nonlinear spectroscopy for the kitaev spin liquid, *Phys. Rev. Lett.* **124**, 117205 (2020).
- [410] Rahul M. Nandkishore, Wonjune Choi, and Yong Baek Kim, Spectroscopic fingerprints of gapped quantum spin liquids, both conventional and fractonic, *Phys. Rev. Research* **3**, 013254 (2021).
- [411] Zi-Long Li, Masaki Oshikawa, and Yuan Wan, Photon echo from lensing of fractional excitations in tomonaga-luttinger spin liquid, (2021), [arXiv:2103.08122 \[cond-mat.str-el\]](https://arxiv.org/abs/2103.08122) .
- [412] S. A. Parameswaran and S. Gopalakrishnan, Asymptotically exact theory for nonlinear spectroscopy of random quantum magnets, *Phys. Rev. Lett.* **125**, 237601 (2020).
- [413] Michele Fava, Sounak Biswas, Sarang Gopalakrishnan, Romain Vasseur, and S. A. Parameswaran, Hydrodynamic non-linear response of interacting integrable systems, (2021), [arXiv:2103.06899 \[cond-mat.str-el\]](https://arxiv.org/abs/2103.06899) .
- [414] Andrew D King, Juan Carrasquilla, Jack Raymond, Isil Ozfidan, Evgeny Andriyash, Andrew Berkley, Mauricio Reis, Trevor Lanting, Richard Harris, Fabio Altomare, *et al.*, Observation of topological phenomena in a programmable lattice of 1,800 qubits, *Nature* **560**, 456–460 (2018).
- [415] R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, M. H. Amin, K. Boothby, P. Bunyk, C. Deng, C. Enderud, S. Huang, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, T. Lanting, R. Li, T. Medina, R. Molavi, R. Neufeld, T. Oh, I. Pavlov, I. Perminov, G. Poulin-Lamarre, C. Rich, A. Smirnov, L. Swenson, N. Tsai, M. Volkmann, J. Whittaker, and J. Yao, Phase transitions in a programmable quantum spin glass simulator, *Science* **361**, 162–165 (2018).

- [416] D wave Systems Inc, D-wave system documentation, <https://docs.dwavesys.com/docs/latest/index.html>.
- [417] Claudio Castelnovo, Mark I. Dykman, Vadim N. Smelyanskiy, Roderich Moessner, and Leonid P. Pryadko, Dephasing with strings attached, *Phys. Rev. B* **97**, 085121 (2018).
- [418] M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*, Applied mathematics series (Dover Publications, 1965).