

Exotic Ground States and Dynamics in Constrained Systems

Dissertation
zur Erlangung des akademischen Grades
Doctor rerum naturalium
(Dr. rer. nat.)

vorgelegt von
Benedikt Andreas Placke

durchgeführt am
Max-Planck-Institut für Physik komplexer Systeme, Dresden

Wissenschaftliche Betreuung durch
Prof. Dr. Roderich Moessner

vorgelegt der
Fakultät Physik
Bereich Mathematik und Naturwissenschaften
Technische Universität Dresden

Eingereicht am 27. Juni 2023

1. Gutachter: Prof. Dr. Roderich Moessner
2. Gutachter: Prof. Dr. Matthias Vojta
3. Gutachter: Prof. Dr. Peter Holdsworth

Verteidigt am 22. August 2023

Acknowledgements

There is a number of people without which this thesis would not have been the same. Most importantly, I want to thank my supervisor Roderich Moessner for sharing with me his wisdom while always encouraging me to explore new ideas.

I also had the pleasure to collaborate with several outstanding scientists, to whom I am deeply grateful. In Dresden, I am particularly thankful to Owen Benton, who was always ready to explain the physics of pyrochlores to me with great patience and in great detail, and Robin Schäfer, with whom I had a great time modeling pyrochlore materials. I would also like to thank Grace Sommers and Shivaji Sondhi for the collaboration on hardcore spin models, Nikolas Breuckmann for a long collaboration on decoding hyperbolic surface codes and for having the staying power it required, and, of course, all my collaborators in experiment, especially Evan Smith, Danielle Yahne, Bruce Gaulin, Stuart Calder, and Kate Ross.

I am also grateful to many people at the Max Planck Institute for the Physics of Complex Systems in Dresden for making my time here enjoyable. This includes in particular Adam McRoberts, Jonathan Nilsson Hallén and Philippe Suchsland, whom I thank for stimulating discussions during our weekly meetings. I would also like to thank the non-academic staff of PKS for making things run smoothly at all times.

For the realization of this document specifically, I acknowledge support and helpful comments from Adam McRoberts, Christos Kourris, Franz-Josef Placke, Jonathan Nilsson Hallén, Owen Benton, Pedro Côtoli, Peng Rao, Philippe Suchsland, and Roderich Moessner

Finally, I would like to thank my family and friends for their continuous support and for making my life as enjoyable as it is.

Contents

1. Introduction	1
1. Simple Models	5
2. Order through Entropy and Slow Dynamics in Constrained Models	7
2.1. Order through Entropy	7
2.1.1. Onsager’s hard rods	8
2.1.2. Frustrated magnets	9
2.2. Jamming and Glassiness	13
2.2.1. Supercooled liquids and the ideal glass	13
2.2.2. Kinetic arrest at zero temperature: jamming	15
2.3. Fractons and Higher-Rank Gauge Theories	17
2.3.1. A simple example: the X-Cube model	18
2.3.2. Tensor gauge theory and higher-moment conservation laws	20
3. Jamming Dynamics of Hardcore Spin Models	23
3.1. A Lattice Model of Hard Spheres: Hardcore Spins	24
3.1.1. From hard spheres to hardcore spins	24
3.1.2. Equilibrium phase diagram on the square lattice	25
3.1.3. Equilibrium phase diagram on the triangular lattice	27
3.2. “Jamming” Hardcore Spins	29
3.2.1. A compression algorithm for hardcore spins	29
3.2.2. Jamming hardcore spins: failure of compression	30
3.3. Phase Ordering Kinetics Near Jamming	32
3.3.1. Square lattice: vanishing vortex mobility	33
3.3.2. Triangular lattice: vanishing domain wall mobility and slow relaxation of local clusters	36
3.4. Conclusion and Outlook	37
4. An Ising Fracton Spin Liquid in Two Dimensions	39
4.1. The Honeycomb model	40
4.2. Excitations and Ground-State Degeneracy	42
4.2.1. Conservation laws and fracton excitations	42
4.2.2. Extensive ground state degeneracy	44
4.3. Monte-Carlo Simulation of Fracton Phases	45
4.3.1. Why local algorithms fail	45
4.3.2. Moving fractons in pairs	46
4.4. Thermodynamics and Correlations	49
4.5. Conclusion and Outlook	51

II. Artificial Systems	53
5. Quantum Error Correction and Decoding	55
5.1. From Classical to Quantum Error Correction	55
5.1.1. Classical repetition code	55
5.1.2. Shor's nine-qubit code	56
5.2. Kitaev's Toric Code	58
5.2.1. Stabilizers and codespace	58
5.2.2. Logical operators	59
5.2.3. Decoding the toric code	61
5.3. The Statistical Mechanics of Decoding	63
5.3.1. Optimal decoding as a statistical mechanics problem	63
5.3.2. Thresholds and phase transitions	67
5.3.3. Beyond the toric code and phenomenological noise	69
6. The Statistical Mechanics of Decoding Hyperbolic Surface Codes	71
6.1. Hyperbolic Surface Codes	72
6.1.1. Hyperbolic surfaces	72
6.1.2. Hyperbolic surface codes	74
6.1.3. The statistical mechanics of decoding hyperbolic surface codes	75
6.2. The Ising Model and its Dual in Hyperbolic Spaces	77
6.2.1. Duality in the hyperbolic Ising model	77
6.2.2. The random-bond Ising model	81
6.2.3. Numerical phase diagram of the RBIM	84
6.2.4. Numerical phase diagram of the dual RBIM and threshold values	87
6.3. Conclusion and Outlook	90
III. Materials	93
7. The Hunt for Quantum Spin Ice	95
7.1. Classical Spin Ice	95
7.1.1. The ice rule and emergent Maxwell theory	96
7.1.2. Spin ice materials	97
7.2. Quantum Spin Ice	98
7.2.1. A quantum liquid	99
7.2.2. Emergent $U(1)$ gauge theory	100
7.2.3. Phase diagram of the XXZ-model	101
7.3. Dipolar-Octupolar Pyrochlore Materials	102
7.3.1. Hierarchy of energy scales	102
7.3.2. Classification of crystal field doublets	103
7.3.3. Exchange Hamiltonian and coupling to external field	105
7.3.4. Phase diagram of dipolar-octupolar pyrochlores	106
8. Ground State Selection in Dipolar Octupolar Pyrochlores at Zero field	111
8.1. Scattering Neutrons off Dipolar-Octupolar Pyrochlores	112
8.1.1. Neutron scattering cross section beyond the dipole approximation	113
8.1.2. Neutron scattering signatures of dipolar-octupolar spin ices	114

8.2.	Case for a $U(1)_\pi$ Quantum Spin Liquid Ground State in the Dipole-Octupole Pyrochlore $Ce_2Zr_2O_7$	116
8.2.1.	Model Hamiltonian and fitting procedure	117
8.2.2.	Results	119
8.3.	Dipolar Spin Ice Regime Proximate to an All-In-All-Out Néel Ground State in the Dipolar-Octupolar Pyrochlore $Ce_2Sn_2O_7$	122
8.3.1.	Specific heat	123
8.3.2.	Magnetic susceptibility	125
8.3.3.	Structure factor	126
8.3.4.	Overall best fit	128
8.4.	Conclusion and Outlook	128
9.	Field-Induced Spin Chains in Dipolar Octupolar Pyrochlores	131
9.1.	Mean-field Phase Diagram in [110] Field	132
9.1.1.	Zero field limit	132
9.1.2.	Polarization of dipolar spin ice and ordered phases	133
9.1.3.	Octupolar chain phase	135
9.2.	Order by Quantum Fluctuations	138
9.2.1.	Effective triangular lattice Ising model	138
9.2.2.	Order by disorder phase diagram	142
9.2.3.	Field tunability	146
9.3.	Experimental Relevance	147
9.3.1.	Disordered dipolar chain state in $Nd_2Zr_2O_7$	147
9.3.2.	Proposed octupolar spin liquids in $Ce_2Zr_2O_7$ and $Ce_2Sn_2O_7$	147
9.4.	Modeling $Ce_2Zr_2O_7$ at Finite Field	148
9.4.1.	(Re-)Fitting specific heat at finite field	148
9.4.2.	Evidence for the Chain _Y * phase?	151
9.4.3.	Integrated structure factor	152
9.4.4.	Dynamics	156
9.5.	Conclusion and Outlook	158
	Appendices	161
A.	Compression protocols for hardcore spins	163
A.1.	Lubachevsky-Stillinger Compression Protocol	163
A.2.	Softcore Protocol	163
B.	Kink-free relaxation of hardcore spins	167
B.1.	Generation of Kink-Free Initial States	167
B.2.	Compression of Kink-Free States	167
C.	High-Temperature Series Expansion of the Hyperbolic Ising model	169
C.1.	Biconnected Graph Expansion of Inverse Susceptibilities	169
C.2.	Analysis of the Series	170
D.	Neutron Scattering Cross Section of Dipolar-Octupolar Pyrochlores	173
D.1.	Matrix Elements of Magnetization Operator	173
D.2.	Structure Factor	175
D.3.	Explicit Expression for Ce^{3+} Doublet	175

E. Semiclassical Monte-Carlo Molecular Dynamics	179
E.1. Landau-Lifshitz Dynamics	179
E.2. Generating the Initial States	179
E.3. Molecular Dynamics Evolution	180
E.4. Spectral Classical-to-Quantum Conversion	181
F. Real Space Perturbation Theory of Chain Phases in Dipolar-Octupolar Pyrochlores	183
F.1. Perturbation Theory	184
F.2. Degeneracy Lifting	185
F.3. Chainy Phase	189
F.3.1. Effective coupling J_{ch}^y	190
F.3.2. Effective coupling $J_{\text{ch}}^{y'}$	192
G. Spin-Wave Theory of Chain Phases in Dipolar-Octupolar Pyrochlores	193
G.1. Details of the Numerical Calculations	194
G.2. Survival of Full Classical Degeneracy when $J_x = 0$ or $J_y = 0$ in the Dipolar Chain Phases	197
G.3. Effect of Finite Temperature	197
G.3.1. Low temperature expansion	198
G.3.2. Free energy of magnons at finite temperature	199

1. Introduction

Ever since its emergence as an independent subfield, condensed matter physics has been a fruitful avenue for theory and experiment to realize novel ideas and phenomena. By nature, physics is an experimental science and condensed matter is of course no exception. Indeed, especially in its early days, progress in the field was driven largely by experimental discoveries. Examples range from the first observation of superconductivity [1], the theoretical explanation of which five decades later [2] might be considered as one of the foundational pillars of modern condensed matter theory. In the following decades, further experimental surprises such as the (fractional) quantum Hall effect [3, 4] and unconventional superconductivity [5] stimulated much theoretical work [6, 7], as did the quest to understand the nature of long established experimental phenomena such as glassiness [8–10].

Of course, the relationship between theory and experiment is not strictly a one-way road where experiments yield data that theorists try to make sense of. As condensed matter theory matured and especially after Anderson famously realized that “more is different” [11], theory has developed such that today, we are able to “discover” entirely new physical phenomena, without direct motivation in any particular experiment [12–18].

The contribution of this thesis

The overarching theme of this thesis is the question of how constraints influence collective behavior. Constraints are crucial in shaping both static and dynamic properties of systems across diverse areas within condensed matter physics and beyond. For example, the simple geometric constraint that hard particles cannot overlap at high density leads to slow dynamics and jamming in glass formers. Constraints also arise effectively at low temperature as a consequence of strong competing interactions in magnetic materials, where they give rise to emergent gauge theories and unconventional magnetic order. Enforcing constraints artificially in turn can be used to protect otherwise fragile quantum information from external noise.

This thesis in particular contains progress on the realization of different unconventional phases of matter in constrained systems. The presentation of individual results is organized by the stage of realization of the respective phase. Novel physical phenomena after conceptualization are often exemplified in simple, heuristic models bearing little resemblance of actual matter, but which are interesting enough to motivate efforts with the final goal of realizing them in some way in the lab. One form of progress is then to devise refined models, which retain a degree of simplification while still realizing the same physics and improving the degree of realism in some direction. Finally, direct efforts in realizing either the original models or some refined version in experiment today are mostly two-fold. One route, having grown in importance rapidly during the last two decades, is via the engineering of artificial systems realizing suitable models. The other, more conventional way is to search for realizations of novel phases in materials.

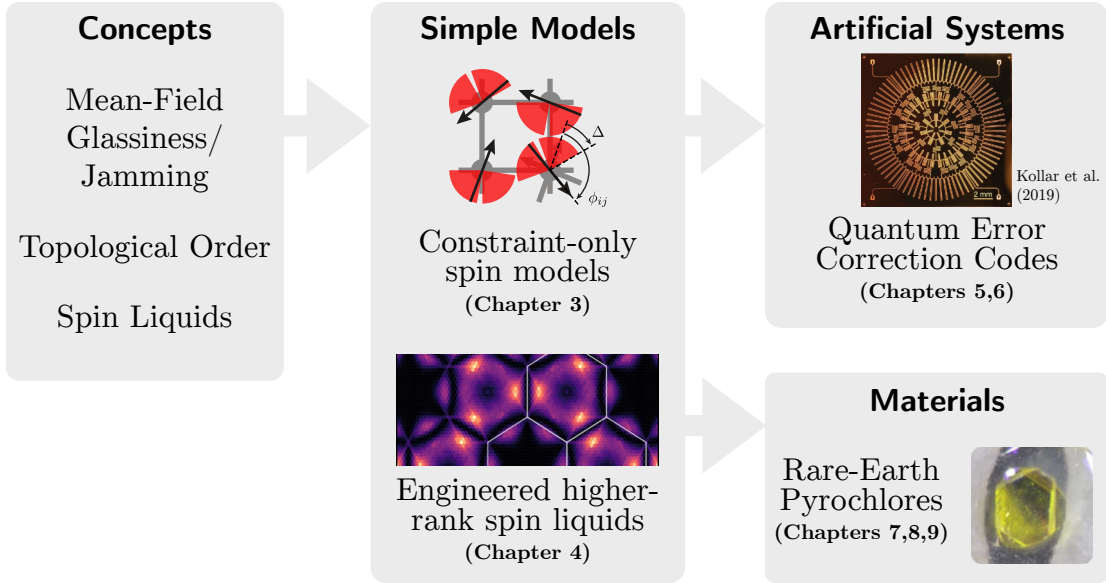


Figure 1.1.: From conceptualization of novel physical phenomena to the lab. Novel theories are often shown to be viable in toy models before eventually being realized in experiment. The latter happens either in artificially engineered systems or material platforms. This thesis contains a broad study of novel phenomena in constraint systems, at different stages of their journey towards realization.

The thesis is divided into three parts, where Part I is devoted to the study of two simple models, while artificial systems and real materials are the subject of Part II and Part III respectively. A pictorial overview of the systems studied is shown in Fig. 1.1. Below, the content of each chapter is summarized in more detail.

Summary of Part I

After an introduction to entropic ordering and slow dynamics in chapter 2, in chapter 3 we present a family of models devised as a lattice analog of hard spheres. These are often studied to explore whether low-dimensional analogues of mean-field glass- and jamming transitions exist, but also serve as the canonical model systems for slow dynamics in granular materials more generally. Arguably the models in this family do not offer a close resemblance of actual granular materials. However, by studying their behavior far from equilibrium, we observe the onset of slow dynamics and a kinetic arrest for which, importantly, we obtain an essentially complete analytical and numerical understanding. Particularly interesting is the fact that this understanding hinges on the (in-)ability to anneal topological defects in the presence of a hardcore constraints, which resonates with some previous proposals for an understanding of the glass transition.

Chapter 4 contains the detailed study of a two-dimensional fracton spin liquid. The model is an Ising system with an energy function designed to give rise to an emergent higher-rank gauge theory at low energy. We show explicitly that the number of zero-energy states in the model scales exponentially with the system size, establishing a finite residual entropy. A purpose-built cluster Monte-Carlo algorithm makes it possible to study the behavior of the model as a function of temperature.

We show evidence for a first order transition from a high-temperature paramagnet to a low-temperature phase where correlations match predictions of a higher-rank coulomb phase.

Summary of Part II

Turning away from heuristic models, chapter 5 contains an introduction to quantum error correction, a scheme where constraints are artificially imposed in a quantum system through measurement and feedback. This is done in order to preserve quantum information in the presence of external noise, and is widely believed to be necessary in order to one day harness the full power of quantum computers. Given a certain error-correcting code as well as a noise model, a particularly interesting quantity is the threshold of the code, that is the critical amount of external noise below which quantum error correction becomes possible. For the toric code under independent bit- and phase-flip noise for example, the threshold is well known to map to the paramagnet to ferromagnet transition of the two-dimensional random-bond Ising model along the Nishimori line. Chapter 6 presents the first generalization of this mapping to a family of codes with finite rate, that is a family where the number of encoded logical qubits grows linearly with the number of physical qubits. In particular, we show that the threshold of hyperbolic surface codes maps to a paramagnet to ferromagnet transition in what we call the “dual” random-bond Ising model on regular tessellations of compact hyperbolic manifolds. This model is related to the usual random-bond Ising model by the Kramers-Wannier duality but distinct from it even on self-dual tessellations. As a corollary, we clarify long-standing issues regarding self-duality of the Ising model in hyperbolic space.

Summary of Part III

The final part of the thesis is devoted to the study of material candidates of quantum spin ice, a three-dimensional quantum spin liquid. The work presented here was done in close collaboration with experiment and focuses on a particular family of materials called dipolar-octupolar pyrochlores, which is introduced in chapter 7. This family of materials is particularly interesting because they might realize novel exotic quantum states such as octupolar spin liquids, while at the same time being described by a relatively simple model Hamiltonian. We present a detailed study of ground state selection in dipolar-octupolar pyrochlore magnets and its signatures as observable in neutron scattering. First, in chapter 8, we present evidence that the two compounds $\text{Ce}_2\text{Zr}_2\text{O}_7$ and $\text{Ce}_2\text{Sn}_2\text{O}_7$ despite their similar chemical composition realize an exotic quantum spin liquid state and an ordered state respectively. Then, chapter 9 contains a study of ground state selection of dipolar-octupolar pyrochlores in a magnetic field. Most importantly, we show that the well-known effective one-dimensional physics – arising when the field is applied along a certain crystallographic axis – is expected to be stable at experimentally relevant temperatures. Finally, we make predictions for neutron scattering in the large-field phase and compare these to measurements on $\text{Ce}_2\text{Zr}_2\text{O}_7$.

Part I.
Simple Models

2. Order through Entropy and Slow Dynamics in Constrained Models

This chapter contains a short, pedagogical introduction to two important phenomena frequently encountered in constrained systems. We begin by introducing the concept of order through entropy. This will prove central to the understanding of the static properties of the constraint-only systems studied in chapter 3, but will also be a frequent item of discussion throughout the rest of the thesis.

After that, we introduce two ways in which slow dynamics arise in constrained systems. The first will be jamming and glassiness as seen in a range of models as well as real systems. The common idealized model system for these phenomena, hard spheres, will be the prototype for the “hardcore spins” studied in chapter 3. Finally, we will introduce fracton models, where slow dynamics arises in the form of restricted mobility of excitations as a consequence of conservation laws in the effective low-energy theory. The demonstration of a new model realizing fracton excitations, together with a novel Monte-Carlo algorithm to simulate its thermodynamic properties, will be the main result of chapter 4.

2.1. Order through Entropy

Many ordering transitions, like the formation of crystal structure in solids or the emergence of ferromagnetic ordering in magnets, occur because of a competition between entropy and energy. In these cases, the system’s configurational entropy is larger in the disordered state, while the ordered state has lower internal energy. At high temperature, where the free energy is dominated by entropy, the system is then disordered. However, as temperature is lowered at some point the energy gain of the ordered state, that is the heat released to the environment during the ordering process, outweighs the loss in configurational entropy and the system orders.

Conceptually quite distinct from this scenario are systems which are either athermal [19], or where the state that minimizes the internal energy has a finite, “residual” entropy. The emergence of any spontaneous order, as measured by a suitable order parameter or correlation function, in such systems can only be a consequence of entropy. Clearly, the increase in order must coincide with an increase in configurational entropy. There are now many known examples of such entropy-driven ordering of both of the aforementioned kinds. Here, we begin with a discussion of the archetypical example of an entropy-driven ordering transition in an athermal system: Onsager’s hard rods [20]. In the second part of the section, we explain how the same mechanism is commonly encountered in magnetic systems, where it often lifts accidental degeneracies present because of competing interactions.




Figure 2.1.: Excluded volume for hard rods with parallel (a) and perpendicular (b) orientation. A hard rod excludes the geometrical center of a parallel rod from a volume of the order $V_{\text{ex}} \sim LD^2$ (indicated by a red dashed line), while it excludes a perpendicularly oriented rod from a volume of order $V_{\text{ex}} \sim L^2D$. Figure adapted from Ref. [19].

2.1.1. Onsager's hard rods

As Onsager argued in his seminal 1949 paper [20], a collection of hard rods must undergo a phase transition from a orientationally disordered phase at low density, to a phase with nematic order at large densities. The key of the argument is a competition between orientational and translational entropy. This can be seen by focussing on long rods $L \gg D$, and considering the excluded volume V_{ex} of two rods as a function of their relative orientation θ . The two limiting cases of this are shown in Fig. 2.1. A hard rod excludes the geometric center of another rod with oriented in parallel from a volume scaling as $V_{\text{ex}}^{(\text{par})} \sim D^2L$. In contrast, the excluded volume for a rod with perpendicular orientation scales as $V_{\text{ex}}^{(\text{perp})} \sim L^2D$ and we find

$$\frac{V_{\text{ex}}^{(\text{par})}}{V_{\text{ex}}^{(\text{perp})}} \sim \frac{D}{L} \quad (2.1)$$

Since a smaller excluded volume implies a larger translational entropy, the above implies that while nematic order naturally decreases orientational entropy, it increases translational entropy.

More concretely, the configurational entropy per rod of a dilute gas of hard rods with density ρ and angular distribution $f(\Omega)$ is given by

$$S/N = \underbrace{\log(\lambda^{-3}\rho^{-1}) - \int f(\Omega) \log f(\Omega) d\Omega}_{\text{orientational}} - \underbrace{\frac{\rho}{2} \int f(\Omega) f(\Omega') V_{\text{ex}}(\Omega, \Omega') d\Omega d\Omega'}_{\text{translational}} \quad (2.2)$$

where the first term is the result for an ideal gas with thermal wavelength λ , the second term is the orientational entropy and the third term is the leading correction to the ideal gas because of hard-core repulsion.

Now, remember that the excluded volume V_{ex} is lower for aligned rods than it is for an orientationally disordered gas. Also, note that the excluded volume contribution to the translational entropy in Eq. (2.2) is proportional to the density. This means that at low densities, the system behaves like an ideal gas and there is hence no nematic order. As density is increased, the excluded volume contribution to the translational entropy becomes more important. Using a variational Ansatz for $f(\Omega)$, Onsager was then able to show that for low densities ρ the stable solution of Eq. (2.2) is isotropic ($f(\Omega) = 1/4\pi$), while above a critical density ρ_c the total entropy is maximized by an anisotropic distribution [that is by a non-constant $f(\Omega)$].

Onsager’s theory of hard rods thus makes transparent the mechanism behind order by entropy, that is the competition between several distinct entropic contributions. In general, as “visible” order is increased some corresponding contribution to the entropy decreases and it must be offset by the increase of some other contribution such that the total entropy increases.

As a final note, the entropy in Eq. (2.2) bears an intriguing similarity to the usual formula for the (Helmholtz) free energy $F = U - TS$ with the density taking the role of an effective (inverse) temperature. While such identifications might seem tenuous a similar observation will become quite explicit in the model of constraint-only spins that is the subject of chapter 3. This system undergoes a Kosterlitz-Thouless transition as a function of a tuning parameter quite similar to that of the density in Onsager’s theory of hard rods. This implies that upon coarse graining, one recovers the same effective long-wavelength theory as for the usual XY model, but with the effective stiffness generated purely by entropy rather than energetics.

2.1.2. Frustrated magnets

The second important category of systems where order through entropy is frequently encountered is frustrated magnets. Order through entropy in this context is often called “order by disorder”, after the first discussion of this mechanism in the context of magnetism in the fully frustrated domino model [21], where fluctuations are induced by structural disorder. Presumably, the name is appealing also as a witty reference to the common (mis-)conception of entropy as a “measure of disorder”¹.

Frustration in the context of (classical) magnetic Hamiltonians means that not all microscopic interaction energies can be minimized simultaneously [23]. This often leads to a large number of available states at low temperature and some systems even have an exponentially large number of states (in the system size) that minimize the classical energy [24]. Since this large degeneracy is not protected by any symmetries, they can be lifted by fluctuations, in which case order is again induced by entropy.

An elementary example

It is insightful to discuss all the aforementioned points using an elementary example. Following lecture notes by Chalker [25], let’s consider a fully connected, antiferro-

¹see e.g. the “Oxford dictionary of Science” [22, page 292]




Figure 2.2.: Parametrization of (a) general configuration of four XY spins by three angles and of (b) the ground state configuration ($\mathbf{L} = 0$) of four spins residing on a fully-connected cluster. Panel (c) shows the soft mode of the collinear ground state, with an energy quartic in ϵ .

magnetic cluster of spins \mathbf{S}_j , $j = 1 \dots q$ with n components and Hamiltonian

$$H = \frac{1}{2} \sum_{i \neq j} \mathbf{S}_i \cdot \mathbf{S}_j = \frac{1}{2} |\mathbf{L}|^2 - \frac{q}{2}, \quad (2.3a)$$

$$\mathbf{L} = \sum_{j=1}^q \mathbf{S}_j. \quad (2.3b)$$

For $q > 2$, the terms in the Hamiltonian cannot be simultaneously minimized but instead the ground state is given by all states for which the total magnetization vanishes, $\mathbf{L} = 0$.

To be more concrete, consider the case of XY spins ($n = 2$) on a tetrahedron ($q = 4$). The general state can be (up to symmetry) parametrized by three angles as shown in Fig. 2.2 (a) and has energy

$$H = \cos(\theta) + \cos(\theta + \phi) + \underbrace{\cos(\theta + \phi + \phi')}_{=\mathbf{S}_2 \cdot \mathbf{S}_1} + \cos(\phi) + \cos(\phi') + \underbrace{\cos(\theta + \phi')}_{=\mathbf{S}_3 \cdot \mathbf{S}_1}. \quad (2.4a)$$

The zero-energy manifold then is described by a single angle θ , as shown in Fig. 2.2 (b). To see how fluctuations favor some ground states over the others we inspect the energy penalty of a particular deformation, $\phi = \phi' = \pi - \theta - \epsilon$, for small ϵ . The energy of such a configuration is

$$H = -2 + [1 - \cos(\theta)] \epsilon^2 + \sin(\theta) \epsilon^2 + \mathcal{O}(\epsilon^4). \quad (2.4b)$$

It is evident from the above that fluctuations around different ground states do not have the same cost. The energy of a perturbed state in general scales quadratically with the deformation ϵ except when deforming a collinear state $\theta = 0$, in which case the energy cost is quartic in ϵ . A sketch of this soft mode is shown in Fig. 2.2 (c). At low temperature this soft mode increases the entropy of the collinear states and will, at least for two-component spins, lead to the selection of that state by thermal fluctuations.

Generally, whether the soft mode will lead to the entropic selection of a particular state depends on the dimension of the zero-energy manifold compared to the soft subspace [25–27].

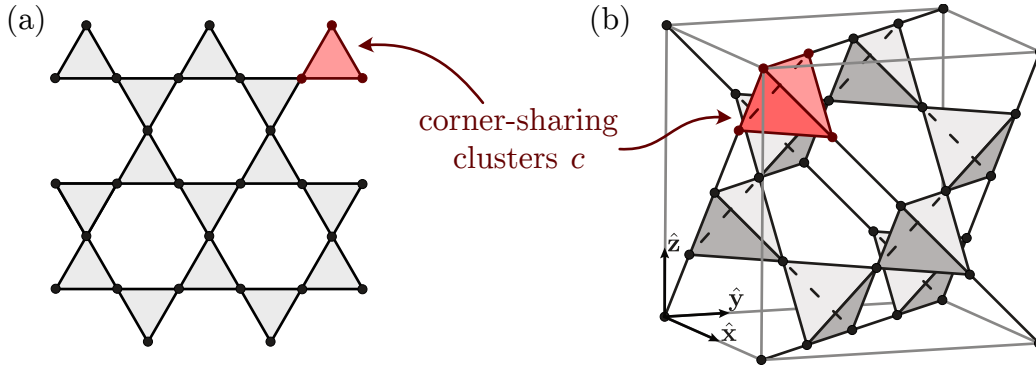


Figure 2.3.: Lattices composed of corner sharing, fully connected clusters. (a) Kagomé lattice, (b) Pyrochlore lattice.

Extended systems

Note that the example presented above, albeit simple, is not merely an intellectual exercise. In fact, there is an abundance of materials where localized magnetic moments form lattices of corner sharing fully-connected clusters. The two most commonly studied such lattices are shown in Fig. 2.3: the Kagomé lattice and the Pyrochlore lattice, consisting of corner-sharing triangles ($q = 3$) and tetrahedra ($q = 4$), respectively. The nearest-neighbor Heisenberg model on these lattices has the important property that its Hamiltonian can be rewritten similarly to Eq. (2.3), as a sum over the local clusters c

$$H_{\text{lat}} = \frac{1}{2} \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{S}_j = \frac{1}{2} \sum_c |\mathbf{L}_c|^2 + \text{const.} \quad (2.5)$$

Since all terms are positive, the ground state of H_{lat} is then given by all states with $|\mathbf{L}_c| = 0 \forall c$.

As mentioned before, many systems with Hamiltonians of the form Eq. (2.5) have a large ground-state degeneracy. This makes possible fascinating physics such as spin liquids, where no long-range order develops even in the limit of the temperature approaching zero $T \rightarrow 0$. Such states will be the subject of study later in this thesis, in particular in chapter 4 and Part III. To establish their existence with certainty however it is extraordinarily important to understand alternative scenarios, which however can also be interesting on their own account

This is all the more important since the temperature scales at which thermal order-by-disorder effects assert themselves, as well as the resulting ordered moment, can be remarkably small. Consider for example the perhaps most intensely studied example of thermal order-by-disorder in an extended (model) system, that is the Heisenberg model on the Kagomé lattice. This model shows a remarkably complex ordering mechanism with a discrete but extensive set of coplanar states being selected by harmonic fluctuations at a temperature around $T \sim 10^{-3}J$ [28–32]. All these states have the same harmonic spectrum [28], but anharmonic fluctuation nevertheless induce long-range order as $T \rightarrow 0$, with an ordered moment about one-order of magnitude below saturation [33].

What about discrete degrees of freedom?

Note that the ordering mechanism described above, while clearly entropic in nature, is of a slightly different flavor than the one encountered in Onsager’s hard rods. There, since the system is athermal, we considered fluctuations *within* the set of configurations satisfying the hard-core constraint and used this to argue that above a certain density the allowed states on average have orientational order. In contrast, in thermal order-by-disorder the ordered state is instead selected by fluctuations which violate the ground state constraint having a finite contribution to the total entropy of the system at finite temperature.

Generally, in frustrated magnets both flavors of order through entropy are possible. However, there is an important distinction in that thermal order-by-disorder as discussed above occurs much more naturally in systems with continuous degrees of freedom, which allow arbitrarily small deformations of the zero-energy states. Conversely, in (classical) systems with discrete degrees of freedom such as Ising models there is a finite energy gap which suppresses excitations above the ground state exponentially at temperatures below the gap. Intuitively, they can hence only lead to a selection of a subset of states if this energetic penalty is set off by multiplicity. To the best of my knowledge, no example exists of a system where an extensive degeneracy is lifted by this mechanism. There are however cases of subextensive degeneracies being lifted by an extensive number of excitations [34]. However, the above constitutes no proof that extensive degeneracies cannot be lifted by thermal fluctuations in systems with discrete degrees of freedom. To either construct such example or refine this line of thought to rule out its existence hence remains a still open question.

Importantly, fluctuations out of the constraint manifold need not be driven by temperature but can, for example, originate from structural disorder [21] or quantum fluctuations [25] instead. Since the latter case is particularly relevant for the work on field-induced one-dimensional physics in pyrochlore materials presented in chapter 9, we discuss it in some more detail in the following.

Quantum order-by-disorder

Naturally, the ground state manifold of a classical spin Hamiltonian need not to be preserved when spin variables are promoted to quantum operators. It is generally ill defined to draw a straightforward between quantum and thermal order by disorder. However, if the classical ground states of the Hamiltonian are not connected by any finite order of perturbation theory, one usually expects that quantum fluctuations merely “dress” the classical states (eg. by renormalizing the value of the ordered moment) rather than changing their nature qualitatively [35].

In this semiclassical limit, zero point fluctuations on top of a classical ground state configuration $\{\mathbf{S}_j\}$ can be treated within the $1/S$ -expansion (see App. G). At leading order, this generates an effective low-energy Hamiltonian of the form

$$H_{\text{LSWT}}[\{\mathbf{S}_j\}] = \sum_{\mathbf{k}} \hbar \omega_{\mathbf{k}}[\{\mathbf{S}_j\}] \left(\hat{n}_{\mathbf{k}} + \frac{1}{2} \right), \quad (2.6)$$

where $\hat{n}_{\mathbf{k}}$ is a Boson number operator and $\omega_{\mathbf{k}}$ is called the spin wave dispersion. As indicated in Eq. (2.6), the spin wave dispersion generally depends on the classical ground state configuration. In particular, in the presence of an accidental degeneracy

as in Eq. (2.3), it is generally different for different ground states and so the leading correction to the classical ground state energy

$$\mathcal{E}_0[\{\mathbf{S}_j\}] = \frac{\hbar}{2} \int_{\text{BZ}} \omega_{\mathbf{k}} d\mathbf{k}, \quad (2.7)$$

will lift such accidental degeneracies.

It is worth noting that while the magnitude of energy splitting induced by fluctuations is often tiny, it can be relevant and even observed experimentally, as demonstrated in the Pyrochlore magnet $\text{Er}_2\text{Ti}_2\text{O}_7$ [36–38].

2.2. Jamming and Glassiness

The study of glasses, jamming and dynamical arrest is a rich subject in theory, but also one of practical relevance for a great variety of industrial applications. To name a particular popular example, the development of the latest generation of Corning’s Gorilla® glass, covering the display of many modern smartphones, benefitted significantly from recent theory advances [39, 40]. More generally, research in the field applies to a wide range of related phenomena in physical systems such as spin glasses [41, 42] and protein folding [43] but also to problems in computer science [44].

This section provides an overview of structural glassiness and jamming. Ultimately, it is the discussion of athermal slow dynamics and random jamming that will be most relevant for the work presented in chapter 3. Still, it is worthwhile to start this section with a brief discussion of supercooled liquids and the idea of an ideal glass transition, which remains central to the field as a whole. After that, we move on to colloidal suspensions which are believed to be well described as a collection of hard particles with only contact interaction. Finally, we discuss some of the theory results on hard spheres, the most common idealized model system considered by theorists studying glassiness and jamming in athermal systems.

2.2.1. Supercooled liquids and the ideal glass

Phenomenology of supercooling

When a liquid is cooled sufficiently fast below its freezing temperature, crystallization can be avoided and the resulting liquid state is called supercooled. Upon further cooling, the microscopic dynamics of the liquid slows down and eventually the liquid appears frozen in a disordered state on any reasonable experimental time scale. In this case, the supercooled liquid is called a glass [45]. This slowdown of microscopic dynamics and the resulting freezing happens dramatically fast, albeit continuously, within a narrow temperature window. This can be used to define a glass transition temperature T_G empirically as the temperature where the viscosity η reaches some (arbitrary but large) threshold. The rapidly diverging viscosity also limits the temperature range accessible in experiment, since below T_G the relaxation time towards equilibrium becomes prohibitively large.

The transition, if defined in this way, is not a sharp thermodynamic phase transition and all thermodynamic quantities behave smoothly. Accordingly, there is no notion of universality close to it and the temperature dependence of many thermodynamic quantities is rather material-dependent. For example, the temperature T_G

as estimated above depends (weakly) on the way the glassy state is prepared [46, 47]. Still, one can divide the family of glass formers roughly into two categories [45], depending on their temperature dependence of the viscosity. The first category are the so called strong glasses, which includes for example the silicates (eg. SiO_2), where the viscosity can be described reasonably well by an Arrhenius law

$$\eta = A \exp\left(\frac{E}{k_B T}\right). \quad (2.8)$$

The second category, which includes most organic glasses, is called fragile glasses and their viscosity is described by the empirical Vogel-Tammann-Fulcher (VTF) equation [48–50]

$$\eta = A \exp\left(\frac{B}{T - T_0}\right), \quad (2.9)$$

with temperature-independent constants A , B . Formulating a theory describing the behavior observed in experiment remains the central open problem in the theory of glasses.

An ideal glass?

One proposed theory of glasses motivated by experiments on supercooled liquids is the existence of a so called “ideal glass”, that is a disordered state of vanishing entropy in the thermodynamic limit. Its existence is rationalized by considering the entropy decrease of the liquid upon cooling. Assuming that the liquid can be treated as an ensemble of metastable states, its total entropy has to reach the crystal entropy eventually because of its greater specific heat compared to the crystal. Specifically, this would happen at the Kauzmann temperature T_K [8, 9] defined by

$$\Delta S_m = \int_{T_K}^{T_m} \frac{\Delta C_p(T)}{T} dT, \quad (2.10)$$

where ΔS_m is the entropy difference between the liquid and the crystal at the melting temperature T_m and $\Delta C_p(T)$ is the specific heat difference between the liquid and the crystal at temperature T .

In the scenario, underpinning the behavior of glasses is a sharp, thermodynamic transition at the Kauzmann temperature, which is the phase transition from the liquid to the ideal glass. Since the glass transition prevents any experimental system to equilibrate close to T_K , there is however little hope of a direct observation of this transition experimentally. Estimating T_K hence can only be done by extrapolating the liquid’s specific heat below $T_G > T_K$. Such estimates are appealing in that they are numerically close to the value of T_0 in the VTF equation for fragile glass formers [9]. Their validity has been questioned by an analysis of low-temperature structural excitations [9], which however only rules out a specific second-order scenario.

The idea of an ideal glass or Kauzmann transition remains popular. In particular, the existence of a first-order ideal-glass transition can be rigorously established in infinite dimensions [51]. However, the fate of this “random first order transition” [52, 53] in finite dimensions is still unclear. Competing proposals include second order [54], as well as topological-defect driven [55] scenarios.

2.2.2. Kinetic arrest at zero temperature: jamming

A wide variety of physical systems can be treated as effectively athermal, including granular materials such as sand, but also foams, and colloidal suspensions [56, 57]. Colloids are particularly appealing since they are believed to be close realizations of one of the paradigmatic idealized model systems used by theorists, that is hard spheres. Because of this, we will describe how slow dynamics arises in these systems in some detail in the following.

Colloidal suspensions

Colloidal suspensions are collections of small, solid particles (nm– μ m scale) in a liquid. Much of the phenomenology observed in supercooled liquids is also observed in such systems, but with the tuning parameter being the particle density rather than temperature.

When appropriately prepared, the particles generally only interact via hardcore contact interactions [58], so that temperature is believed to play a role only in setting the time scale of diffusion, which happens primarily via brownian motion [59]. The relevant parameters characterizing the system are then the packing fraction ϕ , that is the fraction of volume occupied by the particles, and the polydispersity of the particles, that is the distribution of sizes of particles. While polydispersity is important to avoid crystallization [56, 59, 60], physical quantities are usually studied as a function of packing fraction.

The similarity of behavior observed in colloidal suspension and supercooled liquids is perhaps best illustrated by considering the viscosity of the suspension as a function of packing fraction. It can be fitted using the Dolittle equation [61]

$$\eta = A \exp\left(\frac{B\phi}{\phi - \phi_0}\right), \quad (2.11)$$

which is reminiscent of, but different from, the VTF equation describing the viscosity of fragile supercooled liquids as a function of temperature. The numerical value of the parameter ϕ_0 as fitted from experimental data [62] compares well against some theoretical predictions for the position of the so called “random jamming transition” [56]. In the following, we discuss the idea of random jamming in more detail.

Random jamming – a critical point?

Jamming has long history as subject of study both experimentally as well as in theory. Historically, one motivation was the observation that collections of hard spheres upon compression become incompressible much below the close packing fraction — that is they develop a yield stress in a disordered state. In particular, this happens at a density ϕ_{rcp} (≈ 0.84 in 2D and ≈ 0.63 in 3D) which is remarkably reproducible across a variety of systems both in experiment [62, 64] as well as in simulations [56, 63, 65, 66]. Two examples of such a configuration, taken from reference [63] are shown in Fig. 2.4. This led to the proposal of a random close packing density ϕ_{rcp} [67, 68], supposedly the largest density where a stable packing of hard spheres is disordered. This concept is questionable on two grounds. First, the density at which particles jam depends, albeit weak, on the protocol by which they are packed. For example, when hard spherical particles are poured into a container




Figure 2.4.: Two configurations of 27 jammed disks in a square domain, with periodic boundary conditions. The right configuration contains a rattler. The filling fraction in both cases is $\phi \simeq 0.83$. Reprinted from Ref. [63].

one can reach larger packing fraction than ϕ_{rcp} by pouring at a very slow rate as well as shaking the container horizontally [69]. Second, this definition of random close packing is hard to make mathematically precise, since there is a continuous tradeoff between order and density of any stable packing of spheres [70].

Another route to define the jamming transition is via so called inherent structures. For this, one first defines an energy functional which softens the hard sphere constraint, a common choice being

$$V(r_{ij}) = \begin{cases} \frac{\epsilon}{\alpha} \left(1 - \frac{r_{ij}}{\sigma_{ij}}\right)^\alpha & r_{ij} < \sigma_{ij} \\ 0 & r_{ij} \geq \sigma_{ij} \end{cases}, \quad (2.12)$$

where r_{ij} is the distance between particle i and j and $\sigma_{ij} = \sigma_i + \sigma_j$ is the sum of their radii. Note that for $\alpha \rightarrow 0$, the hard-sphere constraint $r_{ij} > \sigma_{ij}$ is recovered. For a system of spheres with such a “softened” constraint, one can define the pressure

$$p = -L^{-d} \sum_{i>j} r_{ij} \frac{dV}{dr_{ij}} \quad (2.13)$$

which is useful to quantify the extent to which the hard-sphere constraint is violated.

The inherent structures are then defined as those configuration reached by gradient descent along the potential energy surface defined by Eq. (2.12) when starting from random initial states.

Remarkably, in Ref. [56] the authors found that these inherent structures change properties as a function of packing fraction ϕ , in a way that is reminiscent of conventional critical points. To this end, for any given initial state the authors define a “point-J” marking the onset of jamming as follows

1. At sufficiently low density, minimization of the total energy using the conjugate gradient method yields an inherent structure of zero potential energy.
2. Compress, that is increase the radii of spheres uniformly by a small fraction, minimizing the total energy after each step.
3. If the minimization after an increment fails to recover a state with zero energy, decompress. That is decrease the radii of spheres uniformly, minimizing the total energy after each step. The procedure is stopped if the pressure p [Eq. (2.13)] reaches zero at ϕ_c .

To verify that the system does indeed not cross an energy barrier in the second step, it is verified that the final configuration is independent on the size of the increment for every initial state. That the final configuration are indeed jammed is verified by computing the dynamical matrix $M_{i\alpha,j\beta}$, defined such that the potential energy difference δV under infinitesimal displacement $\delta r_{i\alpha}$ can be written as

$$\delta V = \frac{1}{2} \sum_{i\alpha,j\beta} \delta r_{i\alpha} M_{i\alpha,j\beta} \delta r_{j\beta}. \quad (2.14)$$

For a jammed packing, the only zero eigenvalues of M should correspond to global symmetries and so called “rattlers”, that is particles which do not overlap with any other particles. A jammed configuration in two dimensions with a rattler is shown in Fig. 2.4. For the analysis presented in the following, all rattlers are removed from the system.

The ensemble of jammed states produced in this way is then studied for a range of exponents α in Eq. (2.12) and in two and three dimensions. In all cases, the ensemble is found to be isostatic, that is the average number of overlaps z that each spheres has with other spheres above jamming $\phi > \phi_c$ approaches $z = 2d$. In particular, the contact number scales as

$$z - z_c = z_0 (\phi - \phi_c)^\beta \quad (2.15)$$

with $\beta \approx 1/2$ and $z_c = 2d$ and z_0 being non-universal. Similar power laws can be found for the shear modulus, viscosity as well as the frequency at which the density of states of vibrational modes is maximal (called the Boson peak frequency) [57].

The above findings make the idea of random jamming as a critical point appealing, but some open questions remain. For example, the above protocol for generating the ensemble of jammed states is not guaranteed to uniformly sample configuration, raising the question to what extent these findings are truly universal. Also, for some time jamming was interpreted to be the zero-temperature limit of the glass transition [71], a view which has since been refuted by numerical studies [66].

In summary, I believe it is fair to say that while much progress has been made in studying the phenomenology of both the glass and the random jamming transition, the question of what might be the mechanism(s) underlying these phenomena still awaits a comprehensive solution.

2.3. Fractons and Higher-Rank Gauge Theories

In this section, we introduce another route to slow dynamics in constraint models that is slightly different from that of geometric frustration in structural glasses as discussed in Sec. 2.2. Instead, an emergent conservation of the dipole moment (in addition to total charge) renders single local excitations immobile. Any local coupling to a bath will then be ineffective in annealing out defects, leading to glassiness.

Such immobile excitations are called fractons [15, 72–81]. In this chapter, we illustrate how they emerge in two slightly different settings. We begin by discussing an exactly solvable model realizing what is called fracton topological order. This is a gapped phase, characterized by a topologically protected ground state degeneracy which scales exponentially in the linear system size. The second part of the section is dedicated to so called higher-rank gauge theories [82–86]. These field

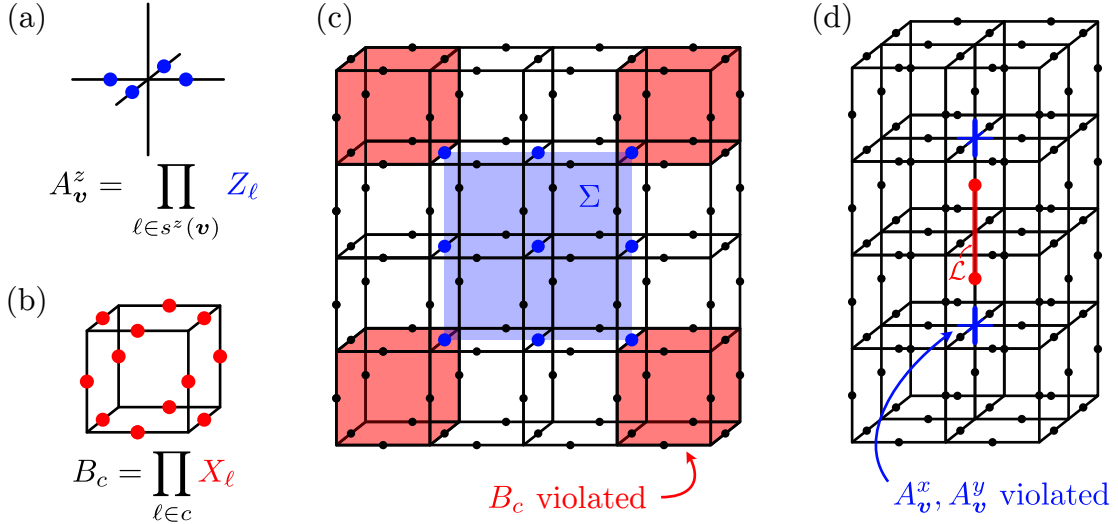


Figure 2.5.: The X-cube model. Its Hamiltonian is the sum over two terms: (a) A_v^α , $\alpha = x, y, z$ is the product of four Pauli- Z operators acting on four coplanar links ℓ incident on the vertex \mathbf{v} (the “X”). The index α denotes the normal direction of the plane that the four edges lie in. (b) B_c is the product of all Pauli- X operators acting on the edges of the cube c . The model’s excitations include: (c) fractons, appearing at the corners of a membrane Σ of Z operators and (d) lineons, appearing at the endpoints of a line \mathcal{L} of X operators.

theories describe a symmetric tensor gauge variable subject to a generalized Gauss law. Importantly, this generalized Gauss law naturally gives rise to dipole moment conservation. This will pave the way to understand the construction of the two-dimensional fractonic spin liquid presented in chapter 4.

2.3.1. A simple example: the X-Cube model

The X-cube model, as first discussed by Vijay, Haah and Fu [77], is defined on the cubic lattice with spin-1/2 degrees of freedoms on the links ℓ of the lattice. The Hamiltonian is the sum of two terms

$$H = -h \sum_{\text{vertices } \mathbf{v}} (A_v^x + A_v^y + A_v^z) - h \sum_{\text{cubes } c} B_c. \quad (2.16a)$$

Here

$$A_v^\alpha = \prod_{\ell \in s^\alpha(\mathbf{v})} Z_\ell \quad (2.16b)$$

is the product of four Pauli- Z operators acting on four coplanar links e incident on the vertex \mathbf{v} . For each vertex, there are three choices of the plane, denoted by its normal direction $\alpha = x, y, z$. This is illustrated in Fig. 2.5 (a). The operators in the second sum are the product of all twelve Pauli- X operators acting on the links of the cube c

$$B_c = \prod_{\ell \in c} X_\ell, \quad (2.16c)$$

which is illustrated in Fig. 2.5 (b).

Using the fact that $X_e, Z_{e'}$ anticommute if $e = e'$ and commute otherwise, it is easy to check that every A_v^α commutes with any B_c as they always share an even number of edges. Hence, all terms in the Hamiltonian can be minimized simultaneously. Since the A_v^α and B operators as products of Pauli operators have eigenvalues ± 1 , the ground state is given by the simultaneous $+1$ eigenstate of all A_v^α and B_c .

Subdimensional particles: Fractons, Planeons, and Lineons

As it is often the case for systems with topological order, it is insightful to consider excitations above the ground state of Eq. (2.16). First, consider the action of a single Pauli- Z operator. Clearly, it commutes with all vertex operators, A_v^α , and anticommutes with four of the cube operators B_c . Generally, acting with Pauli- Z operators on edges within a membrane Σ generates four defects at the corners of this surface, as shown in Fig. 2.5 (c). Because no single pair of such corner defects can be created by local operators, but instead only quadruplets, an isolated defect is completely immobile. A defect with this property is called a fracton.

Evidently, a pair of fractons is mobile since it can be moved by applying a membrane operator such that two corners of the membrane coincide with the two composite fractons. Note however that acting with such membrane can move the pair only within the plane perpendicular to the connection of its two constituents. The pair itself is hence confined to this plane, and because of this property it is called a planeon when viewed as a single dipole-like object.

Lastly, a pair of excitations can also be created by acting on the ground state with a line \mathcal{L} of Pauli- X operators, as illustrated in Fig. 2.5 (d). Two defects are located at the endpoints of the line and they each violate two A_v^μ , with μ the two directions perpendicular to the line. These excitations can be moved by local operators, but their motion is restricted to the direction of the line \mathcal{L} , because line operators in different directions anticommute with different pairs of vertex operators A_v^α . A pair of lineons separated in direction μ is again a planeon with mobility restricted in the plane perpendicular to μ .

Subextensive ground state degeneracy and slow dynamics

The restricted mobility of defects, surely a curious property, comes with physical consequences. First, in a system with periodic boundary conditions, we can reach different ground states by creating, moving, and annihilating pairs of defects non-trivially. Since there are no fully mobile, point-like excitations, moving defects confined to different planes around the system is nonequivalent and the ground state degeneracy (GSD) scales exponentially with the linear system size L . For the X-Cube model it takes the explicit form [77]

$$\log_2 \text{GSD} = 6L - 3. \quad (2.17)$$

The exponential dependence of the GSD on the linear system size is often taken as the characteristic property of fracton topological order [80].

Another important consequence is that the inability of isolated defects to diffuse through the system leads to slow relaxation under local dynamics. In particular, any virtual process has to move *pairs* of defects and is hence exponentially suppressed in the typical defect separation $\xi \sim \exp(h/3T)$ at temperature T , leading to a

characteristic time scale [15]

$$\tau \sim \tau_0 \exp \left[\log \left(\frac{g}{2h} \right) e^{h/3T} \right], \quad (2.18)$$

where $g/2h$ is the microscopic tunneling amplitude, with the energy of the intermediate state $2h$ in the denominator. The time scale in Eq. (2.18) diverges rapidly at low temperature, reminiscent of the behavior of strong glass formers [Eq. (2.8)]. Similar to the situation in model glass formers, the slow relaxation poses a major obstacle in studying fracton phases numerically. In chapter 4, we overcome this difficulty for a specific model by using a hand-crafted cluster Monte-Carlo algorithm that moves pairs of excitations. Slow relaxation may also be advantageous in some cases, for example as a means to robustly store quantum information [74, 79, 87, 88]

2.3.2. Tensor gauge theory and higher-moment conservation laws

The existence of (quasi-)particles with subdimensional mobility might seem rather bizarre. Moreover, while the X-cube model serves as an illustrative and transparent model for their appearance, it is unlikely to be realized in any “natural” setting and even challenging to implement in artificial platforms, due to the necessity to engineer twelve-body interaction of Pauli operators.

Going away from microscopic models, it turns out that fractons arise naturally as a consequence of dipole moment conservation. This is because moving an isolated charge changes the dipole moment. Dipole conservation in turn arises naturally in a simple set of field theories describing a symmetric tensor gauge variable. For concreteness, consider the symmetric tensor version of a Maxwell theory where the gauge sector is invariant under

$$A_{ij} \rightarrow A_{ij} + \partial_i \partial_j \chi. \quad (2.19)$$

The electric field can be defined as the canonical conjugate to A_{ij} such that

$$[A_{ij}, E_{kl}] = i (\delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}) \delta(x - y) \quad (2.20)$$

from which it follows, together with Eq. (2.19), that the electric field obeys a generalization of Gauss’s law [82, 84, 89]

$$\partial_i \partial_j E^{ij} = \rho, \quad (2.21)$$

with $\rho = 0$ for the pure gauge theory.

As in usual electromagnetism, the above Gauss’s law implies the local conservation of charge. Specifically, the total charge in a volume V of space

$$Q = \int_V \rho \, d\mathbf{x} = \int_V \partial_i \partial_j E^{ij} \, d\mathbf{x} = \int_{\partial V} \partial_i E^{ij} \, dn_j \quad (2.22)$$

is given by the local current $\partial_i E^{ij}$ integrated over the boundary of V . This implies the conservation of charge under local operations. Importantly, in the tensor

Maxwell theory one can show an analogous result for the total dipole moment in a volume V

$$P^j = \int_V x^j \rho d\mathbf{x} = \int_V x^j \partial_k \partial_l E^{kl} d\mathbf{x}, \quad (2.23a)$$

$$= \int_{\partial V} x^j \partial_k E^{kl} dn_l - \int_V \partial_k E^{kj} d\mathbf{x} = \int_{\partial V} (x^j \partial_k E^{kl} - E^{jl}) dn_l, \quad (2.23b)$$

where the second line follows from the first via integration by parts. The final step is then analogous to the derivation of charge conservation in Eq. (2.22).

Hence, in the tensor Maxwell theory presented above not only charge is locally conserved but also the dipole moment. Since moving a single charge changes the dipole moment this implies that isolated charges are immobile fractons. Additionally, a local pair of charges forming a dipole \mathbf{p} can move perpendicular to the direction of \mathbf{p} , reproducing the behavior of pairs of defects in the the X-Cube model.

While these similarities are intriguing, note that the X-Cube model and related constructions such as Haah's code [79] are gapped, in contrast to the $U(1)$ gauge theory presented here which has a gapless gauge mode [80]. One way to connect the two constructions is to spontaneously breaking the $U(1)$ symmetry such that fracton topological order is realized as the Higgs phase of a higher-rank gauge theory. Doing this in the simple "scalar charge" theory described above does not preserve subdimensional excitations and one has to consider a modified version to obtain the fracton topological order of the X-cube model upon "Higgsing" [90, 91].

For the purpose of setting the work in chapter 4 in context, the most important conclusion is that higher-rank gauge theories naturally give rise to fractonic excitations with unusual kinetic constraints. Importantly, designing models that realize such higher-rank gauge theories hence provides another route to obtain models that realize fractons as elementary defects. For example, in Ref. [92] a rank-2 $U(1)$ spin liquid was shown to be realized in a classical Heisenberg model on the pyrochlore lattice perturbed by weak Dzyaloshinskii-Moriya interactions. Similarly, in [93] Heisenberg models realizing higher-rank gauge theories are designed systematically via a local constraint similar to the one encountered in Sec. 2.1.2. In both examples, the models do not have fracton excitations due to the continuous nature of the microscopic degrees of freedom. Still, they pave the way towards more physical Hamiltonians that realize the physics relevant to eventually see fractons in the lab. An important step towards this will be presented in chapter 4, were we show that an Ising model obtained via the construction developed in [93] has fracton excitations and we also present strong evidence that the model realizes an emergent higher-rank gauge theory.

3. Jamming Dynamics of Hardcore Spin Models

The work presented in this chapter has appeared as part of the following two articles:

[94] Grace M. Sommers, Benedikt Placke, Roderich Moessner, and S. L. Sondhi. “From hard spheres to hard-core spins”. In: *Phys. Rev. B* 103 (10 Mar. 2021), p. 104407. DOI: 10.1103/PhysRevB.103.104407

[95] Benedikt Placke, Grace M. Sommers, S. L. Sondhi, and Roderich Moessner. “Arresting dynamics in hardcore spin models”. In: *Phys. Rev. B* 107 (18 May 2023), p. L180302. DOI: 10.1103/PhysRevB.107.L180302

Within the realm of condensed-matter physics, examples abound of continuum systems whose critical phenomena can be well captured by simplified lattice models. For instance, the liquid-gas transition can, remarkably, be described using the lattice gas model, which then maps onto the Ising model [96]. Similarly, the Edwards-Anderson model and related spin-glass Hamiltonians provide a fruitful avenue toward understanding random impurities in magnetic alloys [41, 42].

This chapter contains a detailed study of the dynamics of a family of models that we call “hardcore spins”. These models are constructed by analogy to hard spheres (see Sec. 2.2), where the translational degrees of freedom of the spheres are replaced by orientational degrees of freedom of spins on a lattice and the packing fraction as a control parameter is replaced by an exclusion angle. We begin in Sec. 3.1 with a discussion of the equilibrium phase diagram of models on the square and triangular lattice. These exhibit a Kosterlitz-Thouless transition at an exclusion angle Δ_{KT} , with the additional twist that within the critical phase, topological defects can become strictly incompatible with the hardcore constraint.

In Sec. 3.2, we devise compression protocols for hardcore spins and find that *any* protocol that changes the exclusion angle nonadiabatically, if endowed with only local dynamics, fails to compress random initial states beyond an angle $\Delta_{\text{J}} > \Delta_{\text{KT}}$. This coincides with a doubly algebraic divergence of the relaxation time of compressed states towards equilibrium, which is the subject of Sec. 3.3.

We identify a remarkably simple mechanism underpinning this divergent timescale: topological defects involved in the phase ordering kinetics of the system become incompatible with the hardcore spin constraint, leading to a vanishing defect mobility as $\Delta \rightarrow \Delta_{\text{J}}$.

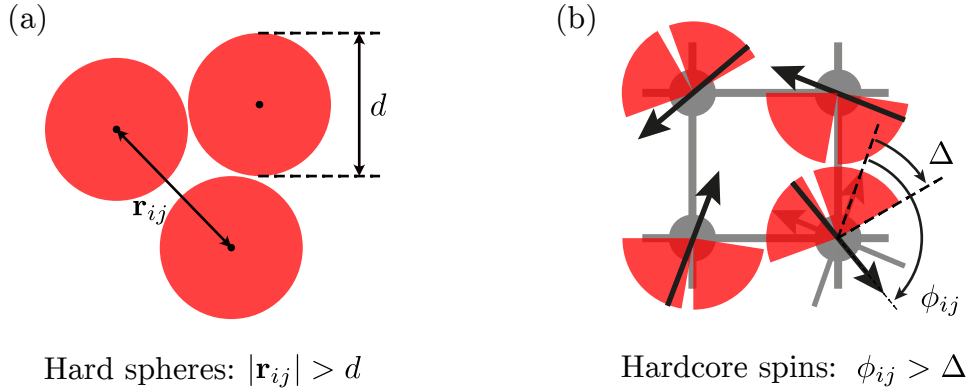


Figure 3.1.: From hard spheres (a) to hardcore spins (b). While for hard spheres, their centers are not allowed to be closer than the sum of their radii, for hardcore spins on a lattice nearest-neighbors must enclose an angle larger than Δ .

3.1. A Lattice Model of Hard Spheres: Hardcore Spins

We begin this chapter by introducing in this section our family of models, which we call “hardcore spins”. We also discuss the equilibrium phase diagram of hardcore spins on the square and triangular lattice.

3.1.1. From hard spheres to hardcore spins

While hard spheres are subject to the non-local constraint that any two sphere centers must not be closer than the sum of their radii, any two hardcore spins that are neighbors on the lattice are not allowed to enclose an angle smaller than the exclusion angle. Explicitly, as illustrated in Fig. 3.1 (a), a system of (monodisperse) hard spheres with diameter d is defined by N position variables $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N$ subject to the constraint

$$|\mathbf{r}_i - \mathbf{r}_j| > d \quad \forall i, j. \quad (3.1)$$

Analogously, we consider vector spins \mathbf{S}_i on a lattice \mathcal{L} that are subject to the constraint

$$\phi_{ij} \equiv \arccos(\mathbf{S}_i \cdot \mathbf{S}_j) > \Delta \quad \forall \langle ij \rangle \in \mathcal{L}. \quad (3.2)$$

where $\langle ij \rangle$ now denotes nearest-neighbor pairs on the lattice, ϕ_{ij} is the enclosed angle between the spins \mathbf{S}_i and \mathbf{S}_j , and we call the parameter $\Delta \leq \pi$ the *exclusion angle*. This is illustrated in Fig. 3.1 (b), where we show a configurations of four XY spins \mathbf{S} on the square lattice. For each spin, we indicate the orientations that are forbidden by the hardcore constraint while keeping all neighbors fixed in red.

The exclusion angle Δ for hardcore spins thus plays the role of the sphere diameter or, equivalently if system size is fixed, the packing fraction in a system of hard spheres. In equilibrium, once the nature of the degrees of freedom and the lattice geometry are fixed, this is the only tuning parameter of the system.

Note that the model Eq. (3.2) separates the states into only two classes – allowed and forbidden – without endowing them either with dynamics or with a notion of (high and low) energy. Thus, any nontrivial correlations are of entropic origin,

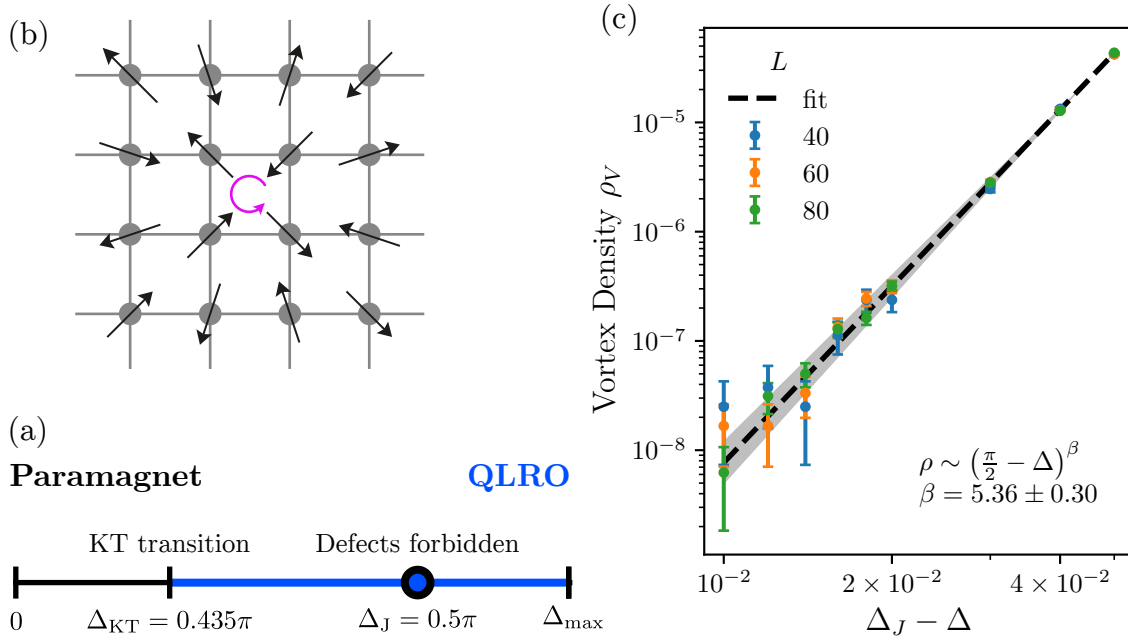


Figure 3.2.: (a) Generic equilibrium phase diagram of hardcore spins, with numerical values as realized on the square lattice. (b) A vortex on the square lattice. The vortex core (indicated in purple), cannot fulfill the hardcore constraint for $\Delta > \Delta_J = \pi/2$. (c) The vortex density vanishes continuously as $\Delta \rightarrow \Delta_J$ from below.

familiar from the nematic ordering of Onsager’s hard rods [20] and its descendants. A general introduction to these systems can be found in chapter 2.

Throughout this chapter, we consider XY spins, parameterized by a single angle, on the square and the triangular lattices.

3.1.2. Equilibrium phase diagram on the square lattice

General expectations and phase diagram

The generic phase diagram of hardcore spins as a function of the exclusion angle Δ , as it is realized on the square lattice, is shown in Fig. 3.2 (a). It is best motivated by considering its two extreme limits. First, at $\Delta = 0$ the system is unconstrained and the system is in a paramagnetic phase, which one expects to extend for some finite range of $\Delta > 0$. Second, on the lattices that we consider the state at Δ_{max} is unique up to global symmetry¹. This state on the square lattice is the Néel state at $\Delta_{max} = \pi$, whereas on the triangular lattice, $\Delta_{max} = 2\pi/3$ and the state is the 120-degree state (sometimes also called $\sqrt{3} \times \sqrt{3}$). These states are thus the analogs of the “close packings” of hard spheres which maximize the packing fraction.

For any $\Delta < \Delta_{max}$, long range order is prohibited by an extension of the Mermin-Wagner theorem to constrained systems [97, 98]. The system instead undergoes a Kosterlitz-Thouless transition at a critical angle Δ_{KT} ($\approx 0.435\pi$ on the square

¹Note that there are lattices on which this is not true: on the Kagomé lattice for example, there is an infinite number of discrete states available that fulfill the hardcore constraint even at $\Delta_{max} = 2\pi/3$. If the lattice is not in some way vertex transitive, the set of states at Δ_{max} can even have finite measure, for example on the pentaheptide lattice.

lattice [99]) into a phase with quasi-long-range (QLR) order, that is algebraically – instead of exponentially – decaying correlations.

Forbidden defects

A detailed discussion of the Kosterlitz-Thouless transition, can be found in Ref. [94] and is not part of this thesis. Instead, we now turn to discuss the role of topological defects. On the square lattice, these take the form of vortices and an example spin texture is shown in Fig. 3.2 (b). Strikingly, and diverging from the simple KT picture, in the hardcore-constrained model topological defects can not only become confined but can even be strictly forbidden by the hardcore constraint. Anticipating the role it will play in the failure of the compression algorithms developed in Sec. 3.2, we call the point at which this happens Δ_J .

This is shown most easily by realizing that on a bipartite graph such as the square lattice, any exclusion model as defined by Eq. (3.2) maps on an inclusion model defined by

$$\tilde{\phi}_{ij} < \Delta_{\text{incl}} \quad \forall \langle ij \rangle \in \mathcal{L}, \quad (3.3)$$

where the spins on one of the two sublattices are flipped

$$\tilde{\mathbf{S}}_j = \begin{cases} \mathbf{S}_j & \text{if } j \text{ in sublattice } A \\ -\mathbf{S}_j & \text{if } j \text{ in sublattice } B, \end{cases} \quad (3.4)$$

so that $\tilde{\phi}_{ij}$ is the enclosed angle between $\tilde{\mathbf{S}}_i$, $\tilde{\mathbf{S}}_j$ and $\Delta_{\text{incl}} = \pi - \Delta$. An antiferromagnetic vortex then maps onto a regular ferromagnetic vortex. Such a vortex however must have a core [94], which is a single plaquette with a winding of 2π . Since there are only four sites around a plaquette, this is only possible if $\Delta_{\text{incl}} > 2\pi/4 = \pi/2$, which implies $\Delta < \pi/2$ in the exclusion model.

Vortices become forbidden at $\Delta_J > \Delta_{\text{KT}}$, that is inside the Kosterlitz-Thouless (KT) phase. In the range $\Delta_{\text{KT}} < \Delta < \Delta_J$, vortices are already confined, but fluctuations lead to a finite density

$$\rho_V = \frac{1}{N_p} \sum_p |V_p| \quad (3.5)$$

where the sum is over all plaquettes p in the lattice, N_p is the number of plaquettes ($N_p = N$ on the square lattice with periodic boundaries), and V_p is the vorticity of the plaquette p

$$V_{p=(i,j,k,l)} = \frac{1}{2\pi} \left[\angle(\tilde{\mathbf{S}}_i, \tilde{\mathbf{S}}_j) + \angle(\tilde{\mathbf{S}}_j, \tilde{\mathbf{S}}_k) + \angle(\tilde{\mathbf{S}}_k, \tilde{\mathbf{S}}_l) + \angle(\tilde{\mathbf{S}}_l, \tilde{\mathbf{S}}_i) \right] \quad (3.6)$$

where $\angle(\tilde{\mathbf{S}}_i, \tilde{\mathbf{S}}_j) \in [-\pi, \pi]$ is the signed phase difference between the spins $\tilde{\mathbf{S}}_i$ and $\tilde{\mathbf{S}}_j$ and the tilde denotes the fact that the spins are taken with respect to the Néel state. The vortex density in equilibrium on the square lattice is shown as a function of exclusion angle Δ in Fig. 3.2 (c). The results are compatible with a power law, that is

$$\rho_V \sim (\Delta_J - \Delta)^\beta \quad (3.7)$$

with $\Delta_J = \pi/2$ and $\beta \approx 5.36 \pm 0.30$.

As we will see later in this chapter, the fact that the vortex density vanishes continuously rather than in a first-order transition has important consequences for the dynamics of the model under compression.

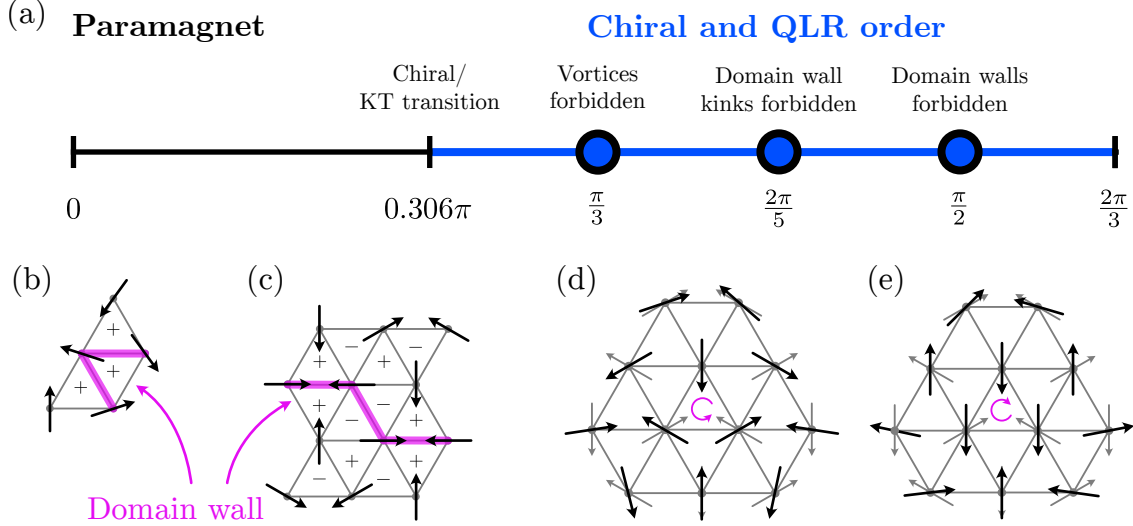


Figure 3.3.: Phase diagram (a) and topological defects (b-e) on the triangular lattice. Because of the XY model on the triangular lattice has an additional chiral symmetry, such that possible defects include domain walls (b,c) in addition to vortices (d,e).

3.1.3. Equilibrium phase diagram on the triangular lattice

Two kinds of defects

The phase diagram of hardcore spins on the triangular lattice is shown in Fig. 3.3 (a). As on the square lattice, there is a low-exclusion-angle paramagnetic and a large-exclusion-angle quasi-long-range ordered phase. However, the situation on the triangular lattice is richer because in addition to the $O(2)$ symmetry of the XY spins, there is a chiral (\mathbb{Z}_2) symmetry associated with the exchange of sublattices. Because of this, there are two kinds of topological defects: domain walls, as well as vortices. The latter are defined with respect to a state of fixed chirality, whose noncollinearity leads to vortices and antivortices having very different free energy costs associated with their creation. To understand this in detail, we must get a better idea of the microscopic spin textures corresponding to different topological defects.

Domain walls: shape matters!

We begin by considering domain walls of chirality. For this, note that at $\Delta_{\max} = 2\pi/3$, not all allowed states can be transformed into each other via global rotations. Instead, there are two incarnations of the “120-degree” or “ $\sqrt{3} \times \sqrt{3}$ ” state, which are related by the exchange of two sublattices. These two states are shown in the top right and bottom left of the right of Fig. 3.3 (c), with the domain wall separating the two chiral domains indicated in pink. To distinguish the two domains quantitatively, we define a vector chirality κ_t on each elementary triangle t of the lattice

$$\kappa_{t=(i,j,k)} = \mathbf{S}_i \times \mathbf{S}_j + \mathbf{S}_j \times \mathbf{S}_k + \mathbf{S}_k \times \mathbf{S}_i. \quad (3.8)$$

In the 120 degree state, the chirality is extremal on all triangles, that is $|\kappa_t| = 3\sqrt{3}/2$, and it has opposite sign on upwards and downwards triangles. This is also indicated for each triangle in Fig. 3.3 (b) and (c).

Observe that domain walls are line-like defects and therefore have an important internal structure, that is, their shape. In particular, it is important to distinguish two kinds. In Fig. 3.3 (d), we show a domain wall that we call “smooth” because it takes no sharp turns or “kinks” as shown in Fig. 3.3 (b). Inspecting the spin texture underlying the smooth domain wall, we recognize the 90-degree structure familiar from the vortex cores on the square lattice. We thus conclude that this kind of domain walls becomes forbidden at $\Delta = \pi/2$. Turning to the domain wall kink, note that such a configuration necessitates three adjacent triangles to have the same sign of the vector chirality κ . This however already implies that the winding number along the loop encircling the three triangles is 2π , which becomes incompatible with the hardcore constraint for $\Delta > 2\pi/5$.

Vortices, again

Turning to vortices, these are defined exactly as on the square lattice, but with respect to a 120-degree state of fixed chirality. In particular we define a vorticity V_t as in Eq. (3.6)

$$V_{t=(i,j,k)} = \frac{1}{2\pi} \left[\angle \left(\tilde{\mathbf{S}}_i, \tilde{\mathbf{S}}_j \right) + \angle \left(\tilde{\mathbf{S}}_j, \tilde{\mathbf{S}}_k \right) + \angle \left(\tilde{\mathbf{S}}_k, \tilde{\mathbf{S}}_i \right) \right] \quad (3.9)$$

where the tilde now denotes the fact that the spins are taken with respect to one of the 120-degree states. Note that this implies that vortices are only well defined in the presence of chiral order. For a uniform choice of chirality, we show a vortex and antivortex in Fig. 3.3 (d) and (e) respectively. In contrast to the square lattice case, the spin textures underlying vortices and antivortices on the triangular lattice are qualitatively different. Still, as we will show in the following, both become forbidden by the hardcore constraint at the same exclusion angle $\Delta = \pi/3$.

First, note that a perfect (meaning unperturbed) antivortex core [the triangle indicated in pink in Fig. 3.3 (e)] is incompatible with the hardcore constraint for any nonzero exclusion angle. What is still possible is a state in which the configuration on the center triangle of the antivortex is perturbed. For this, two of its spins are rotated in opposite directions, away from the third. For the case shown in Fig. 3.3, one possibility of doing this would be to rotate the bottom right spin of the triangle clockwise and the bottom left spin anticlockwise. Via such a rotation by an angle ϵ , we can produce a valid configuration on the center triangle for any $\epsilon < \Delta$. However the vorticity on the central triangle for such state is only equal to -1 as long as $|\epsilon| < \pi/3$, which means that for $|\epsilon| > \pi/3$ it is not an antivortex core. Hence, since an antivortex needs a core its existence is only compatible with the hardcore constraint for $\Delta < \pi/3$.

Second, note that the vortex core as constructed in Fig. 3.3 (d) is exactly the 120-degree configuration of opposite chirality. This might seem to suggest compatibility with the hardcore constraint for all exclusion angles up to Δ_{\max} . However, a vortex is not just a core but an extended object. The next-largest, relevant loop to consider is of length six and takes the form of an upside down triangle centered around the core. The spin configuration of this loop in a perfect vortex configuration [Fig. 3.3 (c)] is clearly only compatible with the hardcore constraint for $\Delta < 2\pi/6 = \pi/3$. At this point, the only spins with any available phase space are those at the three corners. However, any deformation involving these spins already flips the vorticity of the whole loop.

The exclusion angles at which the different topological defects become forbidden by the hardcore constraint are summarized in Fig. 3.3 (a).

3.2. “Jamming” Hardcore Spins

Building on from the discussion of the equilibrium properties of hardcore spins, in this section we turn to study their dynamical properties far from equilibrium.

To this end, in Sec. 3.2.1 we generalize compression algorithms of hard spheres to the case of hardcore spins. Here, “compressing” hardcore spins simply means an increase of the exclusion angle. This is nontrivial because a change in exclusion angle might invalidate the initial state. For this reason, we augment the increase of exclusion angle with dynamics. These dynamics, in a way that is to be made precise, in turn yield a notion of slow and fast compression. We find that while slow compression successfully reaches the unique state at Δ_{\max} [Fig. 3.4 (a)], fast compression fails at $\Delta_J < \Delta_{\max}$ [Fig. 3.4 (b)], that is the compression protocol “jams”².

This can be understood in a surprisingly simple picture: fast compression starting from a random initial state falls out of equilibrium and hence reaches Δ_J with a finite defect density. However for $\Delta > \Delta_J$ defects are strictly incompatible with the constraint and hence the compression protocol fails to increase the exclusion angle beyond that value. Slow compression in contrast, stays in equilibrium. Since the defect density in equilibrium vanishes continuously as $\Delta \rightarrow \Delta_J$ [see Fig. 3.2 (c)], the system reaches Δ_J without any defects and hence does not jam.

3.2.1. A compression algorithm for hardcore spins

The first step towards a jamming of hardcore spins, naturally, is to define a notion of compression, which here will mean an increase of the exclusion angle Δ . Note that this is reminiscent of, but different from, quenches, utilized frequently to study the non-equilibrium dynamics of thermal models. In that case, one induces some dynamics by a sudden change of some system parameter like temperature or a coupling in the Hamiltonian. The issue that sets constrained systems apart is that changing the constraint of the system can make its current state invalid, at which point the behavior of the system is really ill-defined. This makes the development of compression protocols a comparatively nontrivial task.

The problem can be somewhat remedied by simultaneously evolving the state under some dynamics and incrementing the control parameter of the constraint in fixed intervals as much as the current state of the simulation allows. For hard disks, this strategy was implemented by Lubachesvky and Stillinger [63]. In that work, the radii of the disks were increased during a molecular dynamics simulation while keeping the volume fixed. When the compression rate was slow with respect to the time scale set by the molecular dynamics, the final state was approximately close packed. In contrast, when the radii were increased very quickly, the system ended

²Note that when using the word jamming in the context of the hardcore spin model, we refer to the dynamical phenomenology of failed compression. Because of the fundamentally different nature of the degrees of freedom, structural properties of jammed states such as isostaticity do not generalize to the model studied here

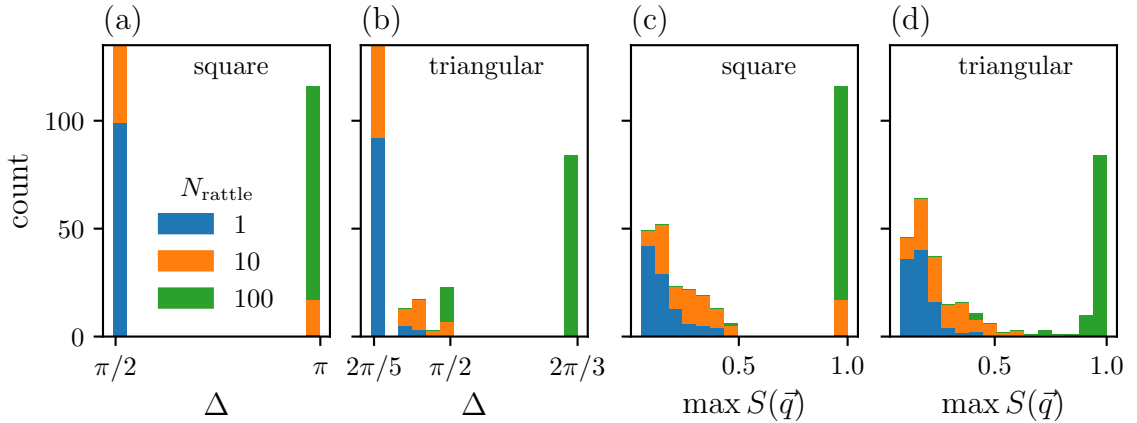


Figure 3.4.: Properties of the final states of the Lubachevsky-Stillinger inspired compression. We study different compression speeds (lower N_{rattle} means faster compression), and the square (a,c) and triangular (b, d) lattice. We show the exclusion angle Δ at which compression fails (a-b) as well as the degree of order (c-d), quantified by the maximum of the structure factor Eq. (3.10).

up jammed in a polycrystalline state (to be contrasted with “maximally random jammed” states [100]), at a density well below close packing.

The Lubachevsky-Stillinger algorithm is straightforwardly adapted to and implemented for hardcore spins, the most subtle point being the choice of dynamics. Arguably the closest analog to the molecular dynamics studied by Lubachevsky and Stillinger is given by continuous-time “Hamiltonian” dynamics. For hardcore spins however, such dynamics are not ergodic since they preserve the local vorticity on each plaquette [94]. Ergodicity can be restored by some nontrivial adjustments and a complementary study to the one presented here can be found in the supplementary material of Ref. [95]. Here however, we focus on a different kind, that is discrete-time, stochastic Monte-Carlo dynamics. Every time step consists of performing a single Monte-Carlo move. Each such move consists of choosing a random site in the lattice, then choosing a random new configuration for the spin on the site and accepting the move if and only if the new configuration is allowed under the hardcore constraint [Eq. (3.2)]. The compression protocol then proceeds as follows. Starting from a random state at $\Delta = 0$, one alternates N_{rattle} Monte-Carlo sweeps (with one sweep consisting of N Monte-Carlo moves) and an increase of the exclusion angle by δ , where the increment δ is chosen as large as possible without invalidating the current state. Fast compression in this context then means large δ/N_{rattle} . The protocol terminates if the increment δ repeatedly falls below a threshold value. This protocol, which we call “hardcore compression” is also given as pseudocode in App. A.

3.2.2. Jamming hardcore spins: failure of compression

Properties of the final states

We run our Lubachevsky-Stillinger inspired hardcore compression protocol 100 times with different values of $N_{\text{rattle}} = 1, 10, 100$ on the square and triangular lattice with $L = 40$ and $L = 42$, respectively, and summarize the results in Fig. 3.4. Analogously




Figure 3.5.: Final states of the Lubachevsky-Stillinger compression protocol on the (a) square (at $\Delta = \pi/2$) and (b) triangular ($\Delta = 2\pi/5$) lattice. In both cases, the rigid backbone (the set of spins which cannot be smoothly rotated while keeping all other spins fixed) is indicated by dashed lines. For the triangular lattice we also highlight the backbone separately on the left.

to what was observed by Lubachevsky and Stillinger, we find strong correlation between the speed of compression set by N_{rattle} , the exclusion angle Δ at which compression fails, and the degree of order in the final states. The latter is quantified by the maximum of the structure factor

$$S(\mathbf{q}) = \frac{1}{N} \sum_{ij} \langle \mathbf{S}_i \cdot \mathbf{S}_j \rangle e^{(\mathbf{r}_i - \mathbf{r}_j) \cdot \mathbf{q}}. \quad (3.10)$$

In particular, for slow compression $N_{\text{rattle}} = 100$, the hardcore compression protocol successfully finds the Néel state at $\Delta_{\text{max}} = \pi$ on the square lattice and the 120-degree state at $\Delta_{\text{max}} = 2\pi/3$ on the triangular lattice. These states are the hardcore-spin analogs of close packing of hard spheres. They are also long range ordered and hence the structure factor takes its maximum at the ordering wave vector, with the maximum value approaching one.

In contrast, for fast compression $N_{\text{rattle}} = 1, 10$ the algorithm generally fails at a much lower exclusion angle. Interestingly, while the distribution is bimodal on the square lattice, on the triangular lattice it is more broad. As will be explained below, this is a direct consequence of the existence of a range of topological defects on the triangular lattice, which become incompatible with the hardcore constraint at different values [cf. Fig. 3.3 (a)].

The degree of order of the compressed states stays well below the equilibrium value. Using the efficient monte-carlo algorithm developed in [94], we find $\max[S(\mathbf{q})] \approx 0.450$ (0.392) for the $L = 40$ ($L = 80$) square lattice at $\Delta = \pi/2$ and $\max[S(\mathbf{q})] \approx 0.384$ on the $L = 42$ triangular lattice. The finite equilibrium value of $\max[S(\mathbf{q})]$ itself is a finite-size effect, and no order is expected in the thermodynamic limit.

Forbidden defects lead to jamming

Importantly, we uncover a remarkably simple mechanism precipitating of what might be called “arresting” or perhaps “jamming dynamics” in our model. In particular, any adiabatic compression protocol, meaning that it stays in equilibrium as Δ is driven through Δ_J , reaches the ordered state with maximal exclusion angle Δ_{max} . In contrast, a nonadiabatic protocol starting from a random initial state at $\Delta = 0$

will fall out of equilibrium and reach $\Delta = \Delta_J$ with nonzero defect density. As the existence of defects is incompatible with the hardcore constraint for $\Delta > \Delta_J$, the protocol will fail beyond this point.

To illustrate that the dynamics here is indeed explained by our aforementioned mechanism, we show in Fig. 3.5 (a) a state at Δ_J for $L = 10$. The “rigid backbone”, that is the set of spins which cannot be perturbed while keeping all other spins fixed, is indicated by dashed black boxes. It is given exactly by the vortex cores, which are completely frozen under local dynamics since the constituent spins enclose an angle of exactly $\pi/2$. Analogously, in Fig. 3.5 (b), we show a jammed state on the triangular lattice at $\Delta = 2\pi/5$. Here, the backbone is again given by topological defects, which now are domain wall kinks. The domain wall kinks are not frozen completely, as there is some available phase space for some spins, shown for clear visibility on the right of Fig. 3.5 (b). However, domain wall kinks are still immobilized sufficiently such that their relaxation time exceeds any reasonable N_{rattle} .

Jamming being facilitated by frozen topological defects is also completely consistent with the results shown in Fig. 3.4. Fast compression on the square lattice jams at $\Delta_J = \pi/2$, where vortices become forbidden. On the triangular lattice the situation is slightly more subtle, but can be understood by considering the same mechanism as on the square lattice: fast compression (that is small N_{rattle}) fails because it arrives with a finite density of topological defects at a point where these become incompatible with Eq. (3.2). Now if there are multiple kinds of defects in the system, these can have different relaxation time scales and hence lead to a more complex dependence of the jamming angle on compression speed.

In this picture, it is somewhat surprising that in Fig. 3.4 (d), there is no peak at $\Delta = \pi/3$, where vortices become forbidden. This is a consequence of the fact that vortices are well defined only in the presence of chiral order, whereas we start from random initial states, that are neither chirally nor QLR ordered and have no well-defined vortex density to begin with. A random initial state will however have a well defined and large domain wall density, which leads to the failure of fast compression (jamming) at $\Delta_J = 2\pi/5$, where kinked domain walls become forbidden. At intermediate speeds, the system is then able to partially equilibrate and reach zero domain-wall-kink density, but still has smooth domain walls leading to jamming at $\Delta_J = \pi/2$. Between these two points, in the range $2\pi/5 < \Delta < \pi/2$, there exist a multitude of local clusters with long but finite relaxation times, leading to jamming of protocols with intermediate compression speeds.

An important, qualitative difference to the square lattice model is that on the triangular lattice the only *vanishing* defect mobility in the model is that of domain walls at $\Delta = \pi/2$. Still, we observe a failure of the compression protocols at other points, where different defects become forbidden without a concomitant freezing. This is not unexpected since a long but finite relaxation time of such defects might still bring compression to a halt. In the language of granular materials, such freezing is *fragile* in that rattling of a jammed state might unjam it.

3.3. Phase Ordering Kinetics Near Jamming

The failure of compression at Δ_J coincides with a doubly algebraic divergence of the relaxation time of compressed states towards equilibrium — as a function of both system size as well as the distance to Δ_J . As is demonstrated in the following, this

can be fully understood in terms of a diffusion-annihilation process of topological defects with a vanishing defect mobility as $\Delta \rightarrow \Delta_J$.

While the LS protocol suffices to demonstrate the presence of jamming dynamics in hardcore spins, it is limited in that it allows only moderate compression speeds, resulting in partial equilibration as evidenced by a finite ordered moment of the final states. Because of this, we implement a second kind of compression protocol, originally developed by Xu et al. [65], which utilizes a softened constraint to compress faster and hence avoid partial equilibration. This protocol is similar to the one described already in Sec. 2.2 and in particular also introduces an energy functional

$$V(\phi_{ij}, \Delta) = \begin{cases} \frac{1}{2} (1 - \phi_{ij}/\Delta)^2 & \text{for } \phi_{ij} < \Delta \\ 0 & \text{for } \phi_{ij} \geq \Delta \end{cases}, \quad (3.11)$$

which is zero if Eq. (3.2) is fulfilled and introduces a quadratic energy penalty if neighboring hardcore spins overlap. The introduction of the soft constraint enables us to use much larger increments δ . This is because one does not have to choose δ such that the current state of the system stays valid, but instead one can choose it such that a conjugate gradient minimization of the total energy after the increment recovers a state with zero energy. We call this protocol “softcore compression”. It is described in detail, including a pseudocode implementation, in App. A.

3.3.1. Square lattice: vanishing vortex mobility

The softcore compression protocol consistently yields far-from-equilibrium states with an ordered moment close to zero. We prepare such states at a range of exclusion angles Δ close to Δ_J and let them evolve under Monte-Carlo time evolution. In Fig. 3.6 (a) and (c) we show the time evolution of the average defect density, ρ_V [Eq. (3.5)], and structure factor at the ordering wave vector, $S(\pi, \pi)$ [Eq. (3.10)], respectively, for a fixed system size of $L = 60$ and a range of exclusion angles. The results are averaged over 1000 initial states. The equilibrium value of the structure factor is computed using the reflect Monte-Carlo algorithm developed in Ref. [94].

As expected from scaling arguments [101], the defect density shows a diffusive decay $\rho_V \sim t^{-1}$ at short times, but at long times $t \sim L^2$ this behavior gives way to an exponential decay with a characteristic relaxation time τ_{defects} . The Néel order parameter $S(\pi, \pi)$ instead decays exponentially for all times, with a characteristic time scale τ_{neel} . Since its finite equilibrium value is an artifact of the finite system size to begin with, we do not expect its decay to be set by scaling.

Both relaxation times as a function of exclusion angle Δ are shown in Fig. 3.6 (b) and (d), for a range of different system sizes. Evidently, the data for both is consistent with a doubly algebraic behavior

$$\tau \sim L^z (\Delta_J - \Delta)^\alpha, \quad (3.12)$$

with $z = 2$ for both, but $\alpha = 3.65 \pm 0.20$ and $\alpha = 3.74 \pm 0.50$ for the relaxation of defect density and structure factor, respectively.

The two power laws in Eq. (3.12) are conceptually quite distinct. The divergence of the relaxation times with system size, $\tau \sim L^z$ owes its existence to universal long-wavelength physics of phase-ordering kinetics under local dynamics and consequently its value $z = 2$ [102] is quite robust. In contrast, to understand the divergence of the

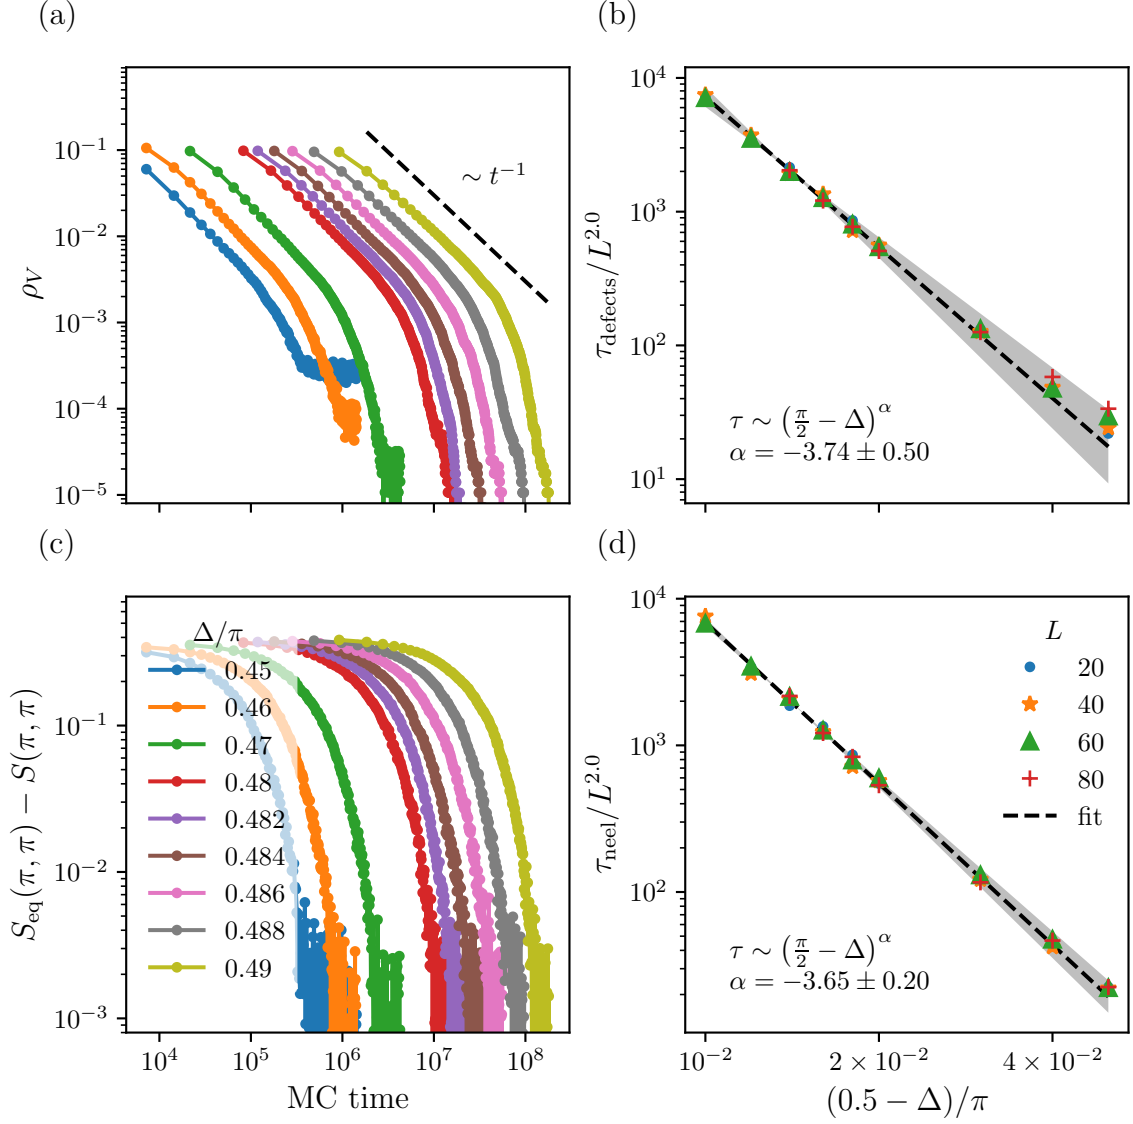


Figure 3.6.: Relaxation of compressed paramagnetic states towards equilibrium. While the density of topological defects (a) initially shows a diffusive decay as expected from scaling arguments, the structure factor (c) decays exponentially towards its equilibrium value for all times. Results are averaged over 10^3 initial states and shown for $L = 60$. The respective relaxation times (b) and (d) diverge with a power law as $\Delta \rightarrow \pi/2$. The area between the two lines with exponent $\alpha \pm \sigma_\alpha$ is indicated in gray.

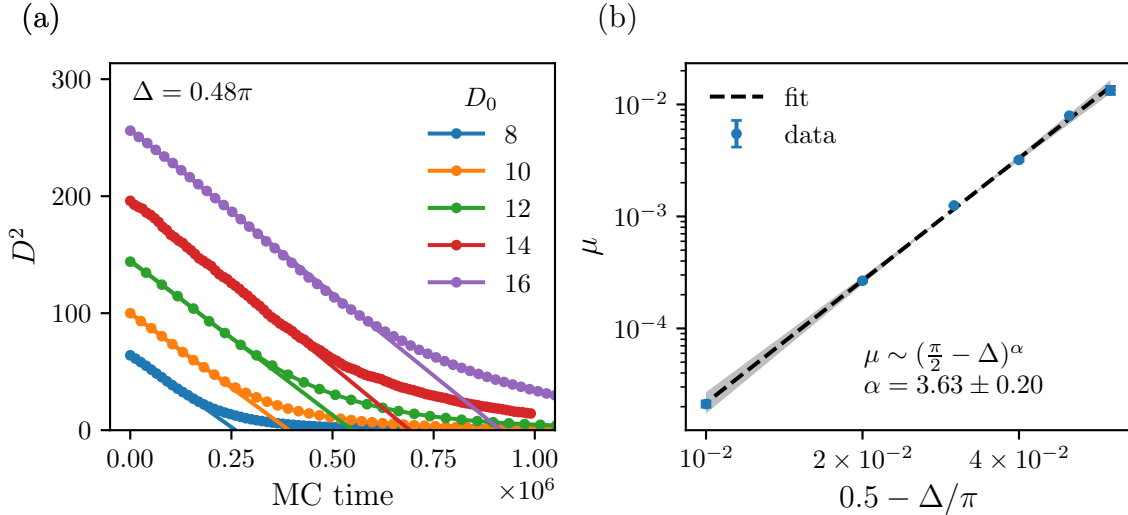


Figure 3.7.: (a) average distance D of a single vortex-antivortex pair prepared on a $L = 60$ square lattice with an initial distance D_0 . Time evolution is performed at exclusion angle $\Delta = 0.48\pi$. For large distances/short times, the time dependence is well described by $D(t) = \sqrt{D_0^2 - \mu t}$. Results are averaged over 10^4 initial states. (b) the mobility μ depends on Δ and vanishes with the same power law as found for the relaxation time of the defect density and the structure factor. The area between the two lines with exponent $\alpha \pm \sigma_\alpha$ is indicated in gray.

relaxation time as the exclusion angle approaches jamming, we have to consider the specific process that drives phase ordering kinetics in the system, that is, diffusion-annihilation of vortex-antivortex pairs.

In the KT phase, a vortex-antivortex pair at distance D is subject to an entropic attractive interaction potential $V(D) \sim \log(D)$, which makes the pair eventually annihilate. For a single pair, initially at distance D_0 , the distance as a function of time then satisfies the differential equation

$$\dot{D}(t) = \mu F[D(t)] = -\frac{\mu}{D(t)} \quad (3.13)$$

where $F(D)$ is the attractive force between the pair and μ is the mobility of vortices. The above equation has a simple solution

$$D(t) = \sqrt{D_0^2 - \mu t}. \quad (3.14)$$

Since the typical initial distance of a vortex-antivortex pair is a property of the initial state, the quantity that sets the time scale of relaxation, and hence the value of the exponent α in Eq. (3.12), must be the mobility μ .

However, this exponent can be readily varied by changing the rules of the dynamical evolution even locally. For example, results from studying phase ordering kinetics under Hamiltonian dynamics plus tunneling are consistent with Eq. (3.12), with $z = 2$ but $\alpha = 1.87 \pm 0.02$

We use the above to compute the mobility as a function of exclusion angle Δ numerically as follows. For each value of Δ , we prepare two isolated vortices at distance D_0 in an otherwise paramagnetic state on a $L = 60$ lattice. This is done by

placing the two vortices into a Néel state and evolving the state using 10^6 Monte-Carlo sweeps, but fixing the vortex and antivortex in place. After this, we evolve the system further under the Monte-Carlo dynamics, but without fixing the vortex-antivortex pair. While doing so, we keep track of the distance of the pair as a function of time. The result of this procedure is shown in Fig. 3.7. In panel (a), we show the distance between the vortex-antivortex pair as a function of Monte-Carlo time, averaged over 10^4 initial states, for a range of initial distances at $\Delta = 0.48\pi$. As expected, for short times and long distances, it obeys the functional form of Eq. (3.14). The mobility can then be extracted from a linear fit to $D^2(t)$. As shown in Fig. 3.7 (b), it vanishes as a power law as $\Delta \rightarrow \Delta_J$ with, for a given local dynamics, roughly the same exponent α as the relaxation times [Fig. 3.6]. However, this exponent can be readily varied by changing the rules of the dynamical evolution even locally. For example, results from studying phase ordering kinetics under Hamiltonian dynamics plus tunneling [95] are consistent with Eq. (3.12), with $z = 2$ but $\alpha = 1.87 \pm 0.02$.

3.3.2. Triangular lattice: vanishing domain wall mobility and slow relaxation of local clusters

Even more so than its equilibrium properties [cf. Fig. 3.3], relaxation dynamics of the model on the triangular lattice are richer but also more complicated. First, because of the higher connectivity of the lattice, the degrees of freedom are generally more constrained and have thus slower dynamics compared to the case of the square lattice at the same exclusion angle. This makes the physics of local clusters with large but finite relaxation times more dominant, which in turn complicates numerical studies. Second, the presence of two kinds of topological defects also leads to additional qualitative features, the detailed study of which is beyond the scope of this work. However, based on general arguments taken from the discussion in Sec. 3.1.3, together with some intuition carried over from the relaxation dynamics of the model on the square lattice and numerical evidence, in the following we still attempt to develop a qualitative picture.

Slow, but finite relaxation beyond the failure of compression

As already discussed in Sec. 3.2.2, a compressed random initial state is not well described in terms of vortices on top of a 120 degree state, but instead by domain walls on top of a chirally ordered state. Accordingly, the short time relaxation of such a state towards equilibrium is governed by the diffusive growth of domains rather than a diffusion-annihilation process of vortex-antivortex pairs.

As shown in Sec. 3.1.3, the first domain-wall shape to become forbidden are “kinks” [Fig. 3.3 (b)] at $\Delta = 2\pi/5$ and domain walls become completely forbidden at $\Delta = \pi/2$. Because the compression protocols of Sec. 3.2 already fail at $\Delta = 2\pi/5$, the range $2\pi/5 < \Delta < \pi/2$ is hard to probe numerically. However based on general arguments we think that phase ordering kinetics as a function of the exclusion angle change as follows. As $\Delta \rightarrow 2\pi/5$, relaxation slows down dramatically, leading to a failure of the compression protocols, and the relaxation times reach a large but finite value. In the range $2\pi/5 < \Delta < \pi/2$, relaxation is slow but does happen. As $\Delta \rightarrow \pi/2$, the relaxation time towards equilibrium diverges.

This picture emerges by considering the microscopic defect mobility, which we showed on the square lattice sets the time scale of relaxation. Observe that while spins participating in domain wall kinks [Fig. 3.3 (b)] have some available phase space, even for $\Delta = 2\pi/5$, this is not true for spins in a smooth domain wall at $\Delta = \pi/2$ [Fig. 3.3 (c)]. This implies that while the mobility of domain wall kinks is small but finite even at $\Delta = 2\pi/5$, it vanishes for smooth domain walls as $\Delta \rightarrow \pi/2$.

The conjectured behavior of relaxation dynamics in the range $2\pi/5 < \Delta < \pi/2$ is hard to verify in numerical simulations. First, the larger connectivity of the triangular lattice which makes the model there generally more constraint which in turn leads to very slow dynamics under local Monte-Carlo updates. Second, in the whole range $2\pi/5 < \Delta < \pi/2$ a multitude of local clusters exhibit long but finite relaxation times, all of which are hard to disentangle from a truly diverging time scale numerically.

To gain access to the range $\Delta > 2\pi/5$ at all, in App. B we subject the system to the additional constraint that kinks in domain walls are strictly forbidden. In this case, we observe a broad distribution of exclusion angles between $2\pi/5 < \Delta < \pi/2$ at which individual runs of the softcore compression protocol (algorithm 2 in App. A) fail. The only instances in which a run terminates outside of this range is if it melts all defects, at which point compression terminates at $\Delta_{\max} = 2\pi/3$. We interpret this as confirmation of the picture drawn above.

3.4. Conclusion and Outlook

In summary, we have provided a detailed phenomenology and comprehensive understanding of arresting dynamics in hardcore spin models, uncovering an intricate interplay of lattice geometry, ordering and defects, and the dynamics at long and short wavelengths. Particularly noteworthy from a conceptual perspective is the role played by the (in)ability to anneal defects under a purely local dynamics. This fundamentally accounts for all phenomenology in this *lattice* model, with considerable added richness as a result of the variability and multiplicity of defects and their configurations. In addition to lending this model intrinsic interest on its own, the central role of defects to the (slow) dynamics resonates with some theories proposed to capture the glass transition of hard spheres [55].

An open question is the applicability of our findings to the random jamming transition of hard spheres. While hardcore spins do not reproduce structural aspects of jammed sphere packings (e.g. isostaticity), the dynamical phenomenology which originally motivated much of the contemporary interest in the topic is reproduced quite well.

4. An Ising Fracton Spin Liquid in Two Dimensions

The work presented in this chapter has appeared as part of the following article:

[103] Benedikt Placke, Owen Benton, and Roderich Moessner. *Ising Fracton Spin Liquid on the Honeycomb Lattice*. arXiv:2306.13151. 2023. DOI: 10.48550/arXiv.2306.13151

Fractons are quasiparticles which are incapable of independent motion. Since equilibration, as seen for example in the last chapter, often happens through the diffusion of quasiparticles, the presence of fractons provides yet another route to anomalous relaxation dynamics and loss of ergodicity. The existence of fractons is well established theoretically [15, 72–81]. They are now understood to arise quite naturally as a consequence of conservation laws present in so called higher-rank gauge theories [82–86], where not only charge is conserved locally but also higher moments (e.g. dipole moment) of the charge density. It is this higher-moment conservation laws that render fractons immobile. The known microscopic theory models that realize higher-moment conservation laws and fracton physics [92, 104–112] are often quite unlikely to arise in materials and even difficult to implement artificially because they involve few-body interactions of high weight. For example, the X-Cube model [77] requires implementing a 12-body interaction term. It is therefore highly desirable to design microscopic models with lower-weight interactions that still possess fracton excitations.

One possibility of making progress in that direction is to focus on classical models, which are more transparent beyond the exactly solvable limit. Classical models that realize higher-rank gauge theories have been constructed [93, 113, 114] so far only with continuous degrees of freedom. The models therefore do not have discrete excitations and hence also no well-defined notion of fracton quasiparticles. In this chapter, we study of a classical Ising model on the honeycomb lattice with local two-body interactions and we present strong evidence that at low temperature the model realizes a higher-rank coulomb liquid [82, 84, 92, 93]. In this Ising model, fractons are naturally present as discrete excitations.

The rest of the chapter is organized as follows. After introducing the model and its mapping to a higher-rank gauge theory in Sec. 4.1, we show explicitly in Sec. 4.2 that defects above the ground state are (type-I) fractons, appearing at the corners of membranes of spin flips. We also show that these membranes can be *locally* combined such that no defects are created. This establishes a number of zero-energy states exponentially large in the number of lattice sites and yields a lower bound on the observed entropy. To study the properties of the model also as a function of temperature, we devise in Sec. 4.3 a novel cluster Monte-Carlo algorithm that

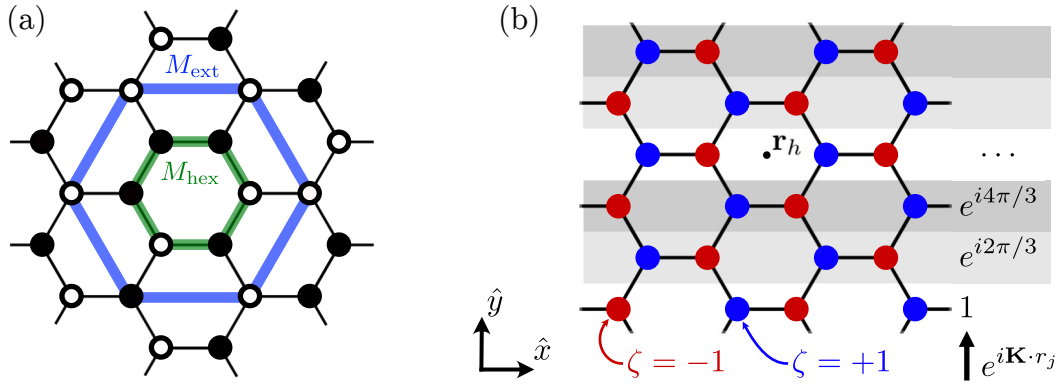


Figure 4.1.: The honeycomb model. (a) Illustration of the constraint M_h with the hexagon and its exterior indicated in green and blue respectively. (b) Sketch of the mapping between the microscopic degrees of freedom σ_j in Eq. (4.1) and the coarse grained field $m_{\mu\nu}$ [Eq. (4.7)].

is purpose-built for the simulation of (type-I) fractonic Ising models. By moving pairs of defects, the algorithm gives access to low temperature properties of large system sizes. In Sec. 4.4, we simulate the model and show evidence for a first-order transition from a high-temperature paramagnet to a low-temperature phase which lacks any sign of conventional order. Instead, correlations precisely match those predicted for a higher-rank Coulomb phase [115].

4.1. The Honeycomb model

We consider a model of classical Ising degrees of freedom σ_j on the Honeycomb lattice, first proposed in Ref. [93] for classical vector spins, with Hamiltonian

$$H = \frac{J}{2} \sum_{\text{hexagons } h} M_h^2 \quad (4.1a)$$

$$M_h = \sum_{j \in h} \sigma_j + \frac{1}{2} \sum_{j \in \langle h \rangle} \sigma_j \quad (4.1b)$$

where the sum in the Hamiltonian is over all hexagons h . For $J > 0$ the ground states of H fulfill the constraint $M_h = 0$ for all hexagons. In the constraint, the first sum is over all six spins in the hexagon h and the second sum is over the exterior of the hexagon $\langle h \rangle$. This is illustrated in Fig. 4.1 (a).

Mapping to higher-rank gauge theory

In Ref. [93], Benton and Moessner showed that the constraint in Eq. (4.1b) can be rewritten as a generalized Gauss's law $\partial_\mu \partial_\nu m^{\mu\nu} = 0$ (see also Eq. (2.21) in Sec. 2.3) of a suitably defined field m .

The construction is as follows. Define real variables t_j, v_j

$$t_j + i\zeta_j v_j = e^{i\mathbf{K}\cdot\mathbf{r}_j} \sigma_j \quad (4.2)$$

where \mathbf{r}_j is the position of site j , $\zeta_j = \pm 1$ is a sign factor that distinguishes between the two sublattices of the honeycomb lattice and

$$\mathbf{K} = \left(0, \frac{4\pi}{3\sqrt{3}a} \right) \quad (4.3)$$

with a being the distance between nearest-neighbor sites. This is illustrated in Fig. 4.1 (b).

Next, we introduce coarse grained, continuous fields $t(\mathbf{r})$ and $v(\mathbf{r})$, which coincide with the microscopic variables at the lattice sites [$t(\mathbf{r}_j) = t_j$ and $v(\mathbf{r}_j) = v_j$]. Assuming that the continuous fields vary slowly with respect to the lattice spacing, we expand them around hexagon centers \mathbf{r}_h to obtain an expression in terms of the microscopic variables t_j, v_j

$$t_j = t(\mathbf{r}_h) + (r_j^\mu - r_h^\mu) \partial_\mu t \Big|_{\mathbf{r}=\mathbf{r}_h} + (r_j^\nu - r_h^\nu) \partial_\nu t \Big|_{\mathbf{r}=\mathbf{r}_h}, \quad (4.4)$$

$$v_j = v(\mathbf{r}_h) + (r_j^\mu - r_h^\mu) \partial_\mu v \Big|_{\mathbf{r}=\mathbf{r}_h} + (r_j^\nu - r_h^\nu) \partial_\nu v \Big|_{\mathbf{r}=\mathbf{r}_h}, \quad (4.5)$$

for a site j that is part of the hexagon h . Inserting the above into the ground-state constraint $M_h = 0$ then yields

$$-\partial_x t + \partial_y t + 2\partial_x \partial_y t = 0. \quad (4.6)$$

Upon defining a traceless, symmetric matrix

$$m = \begin{pmatrix} -t & v \\ v & t \end{pmatrix}, \quad (4.7)$$

Eq. (4.6) becomes

$$\partial_\mu \partial_\nu m^{\mu\nu} = 0, \quad (4.8)$$

which is a generalized Gauss's law of exactly the form discussed in Sec. 2.3.

In Ref. [93], the authors explicitly show that a higher-rank Coulomb liquid is realized at low temperature, by a system of classical three-component spins subjected to a constraint equivalent to Eq. (4.1b). This means that the emergent higher-order gauge theory is in its deconfined phase, the hallmark signature of which is the presence of four-fold pinchpoint singularities in the structure factor [92, 115]. Since the models considered in Ref. [93], as well as related constructions [113, 114], all have continuous degrees of freedom, they have no well-defined single-particle excitations. Hence, it is not possible to isolate and study discrete fractons even in the presence of an emergent higher-rank gauge theory.

While the mapping between the microscopic model and the coarse-grained field $m^{\mu\nu}$ in principle generalizes to the Ising model, it is an open question whether the system still realizes a higher-rank Coulomb phase or whether restricting the degrees of freedom to be discrete yields a set of ground states the average over which no longer corresponds to the deconfined phase of the gauge theory. For the case of the 'conventional' Coulomb liquid, cases are known where the hard-spin Ising and Heisenberg behaviors are (e.g. pyrochlore [116]), and are not (e.g. kagome [33, 117–119]), in accord with that of the soft-spin theory. In the remainder of this chapter, we present evidence for the former case, that is that the Ising model [Eq. (4.1)] does also realize a higher-rank Coulomb phase. Importantly, in this case fractons are naturally present as discrete excitations.




Figure 4.2.: Excitations of the honeycomb lattice model. (a) A single spin flip with respect to the ground state (indicated by a red circle) creates six defects, that is violations of the constraint $M_h = 0$. These are three single ($M_h = \pm 1$) and three double ($M_h = \pm 2$) defects. (b) Flipping an antiferromagnetic bond creates six single defects. Flipping also the next bond along the bond direction (encircled in dashed, orange) moves three of these defects as indicated by the dashed orange arrows. (c) Flipping two bond across a hexagon creates four single defects, which is the lowest-energy excitation possible. Flipping also the next-next hexagon, now perpendicular to the bond direction, moves two defects as indicated again in orange.

4.2. Excitations and Ground-State Degeneracy

4.2.1. Conservation laws and fracton excitations

To understand the hierarchy of excitations of Eq. (4.1), it is instructive to consider first the effect of a single spin flip, as shown in Fig. 4.2 (a). The flipped spin is indicated by a red circle. This results in the constraint $M_h = 0$ being violated on six hexagons. The spin is part of the interior of three hexagons, so that on these $M_h = 2$. We call this a double defect and indicate it by a blue upward pentagon. The flipped spin is also part of the exterior of three hexagons, so that on these $M_h = 1$. We call this a single defect and it is indicated by a blue upward triangle.

While the value of M_h is more transparent when considering the microscopic theory, conservation laws are easier described in terms of the charge ρ of the coarse-grained theory. Evidently, the construction of the field m as summarized in Fig. 3.1 (b) breaks translational lattice symmetry between hexagons, instead distinguishing three sublattices. To relate the constraint M_h to the charge ρ , we have to (arbitrarily) choose one subset of hexagons. One possible choice is indicated in Fig. 4.2 by a darker shade. Denoting this subset as $+$ hexagons and the rest as $-$, the charge is then defined as

$$\rho_h = \begin{cases} M_h & h \text{ is } + \text{ hexagon} \\ -\frac{1}{2}M_h & h \text{ is } - \text{ hexagon} \end{cases}. \quad (4.9)$$

Note that this is completely analogous to the choice of “up” and “down” tetrahedra in the pyrochlore lattice when defining the monopole charge of spin ice [120]. The main difference is that upon Fourier transforming the constraint $M_h = 0 \forall h$, it becomes singular at the Brillouin-zone corners, as opposed to the spin ice constraint which becomes singular at the zone center [93].

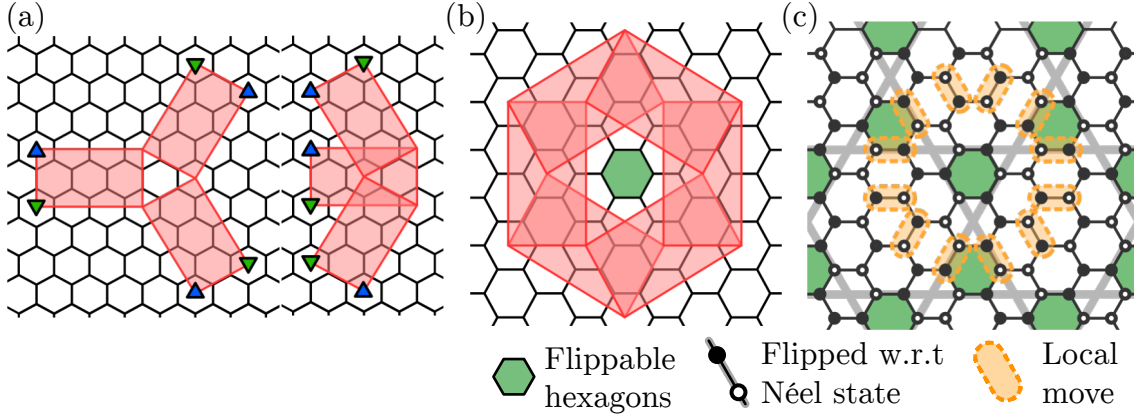


Figure 4.3.: Demonstration of extensive ground state degeneracy. Because of six-fold rotation symmetry, lineons can split into two (a). This can happen in the forward (left) but crucially also into the backward direction (right). The latter allows to close the worldline of lineons, resulting in local moves. The minimal such move, consisting of six interleaved membranes is shown in (b). It is possible to construct a state with an extensive number of such flippable minimal local motifs, shown in (c). The motif in (b) could be centered around any hexagon in (c), implying an extensive ground state degeneracy of the honeycomb model.

Using the definition in Eq. (4.9), it is easy to show that the single spin flip [Fig. 4.2 (a)] conserves total charge $\rho = \sum_h \rho_h$. It also conserves both dipole moments $d_\alpha = \sum_h r_h^\alpha \rho_h$, $\alpha = x, y$, where \mathbf{r}_h is the position of the hexagon h . The first moment of the charge distribution that is not conserved is the quadrupole moment $q_{\alpha\beta} = \sum_h r_h^\alpha r_h^\beta \rho_h$. Any local move is a combination of single flips, and single charges hop by creating dipoles. Conservation of the dipole moment by single spin flips already implies that single charges, corresponding to single defects $M_h = \pm 1$, are fractons.

The effect of flipping two antiferromagnetically aligned Ising spins on a link is shown in Fig. 4.2 (b). The move again creates six fractons. These can also be understood as two lineons, since flipping also the next bond along the bond direction moves three of them as indicated by the dashed orange arrow in Fig. 4.2 (b). Finally, in Fig. 4.2 (c), we show the minimal-energy excitation which creates four fractons. In terms of the generated charge distribution ρ_h , this configuration corresponds to a single quadrupole. Flipping also the next-next hexagon, as indicated in dashed-orange moves a pair of fractons corresponding to a dipole. Generally, single defects appear at the corners of a membrane of flipped hexagons, analogous to the situation in the X-Cube model discussed in Sec. 2.3. This, along with the ability to bind fractons into lineons implies that our model is a “type-I” fracton model, contrasting with “type-II” models where fractons appear on fractal structures and there are no mobile bound states [78].

4.2.2. Extensive ground state degeneracy

Subextensive degeneracy and global moves

The move shown in Fig. 4.2 (b) already implies a subextensive ground state degeneracy of Eq. (4.1), scaling exponentially with the linear system size L . To see this, consider a system with periodic boundary conditions. Starting from the Néel state, which has $E = 0$ and hence is a ground state, we can flip all edges along a line spanning the entire system and by that get to a different state that fulfills the ground state constraint $M_h = 0 \forall h$. One can think of this as creating a pair of lineons as discussed above, moving one of them around the system in a homologically nontrivial way and then annihilating it with its partner again. Since this can be done independently at least for parallel lines, it already implies a ground state degeneracy of at least order $\mathcal{O}(2^L)$, where L is the linear system size.

Extensive degeneracy from local moves

The conventional way of constructing an exponentially large number of ground states in the presence of fractionalized excitations is to create a defect pair on top of a possible ground state and then move and annihilate the pair locally such that a new ground state is obtained. For this, the total path of the pair has to form a closed loop. In spin ice for example, one can create a pair of monopoles and move one of them around around a hexagon to annihilate it with its partner. This yields another ground state, which has all spins around the hexagon flipped with respect to the starting configuration [121, 122]. Since an extensive number of hexagons can be flipped independently, this establishes that the ground state degeneracy has to be exponentially large in the number of sites N .

For fracton excitations, it might seem like the equivalent procedure is only possible in three dimensions. This is because (type-I) fractons appear at the corners of a flipped rectangular membrane. In three dimensions, such a membrane can be folded into a cylindrical shape locally. In two dimensions however, this is not possible. Instead, one has to wrap the rectangle around the entire system, either to the boundary or around the periodic boundary conditions, rendering this a global, rather than a local move. Another possibility is to combine multiple overlapping rectangles to annihilate all corner defects. As I will show in the following, this can be used to construct a local move for hexagonal lattice symmetry lattice which for rectangular lattice symmetry does not exist.

Pairs of fractons that form a lineons can only move perpendicular to the direction of their dipole moment. In the presence of six-fold rotational symmetry however, they can also split into two pairs of lineons moving along the two other directions. Crucially, there are two possibilities conserving total dipole moment. A lineon as shown in Fig. 4.3 (a) can split forward (left) and also backward (right). The latter makes it possible to close the worldline of these defects locally, with the resulting minimal local move shown in Fig. 4.2 (b) and (c). This local move consists of six membranes whose bulk does not overlap, but each corner overlap with exactly one other corner such that no defects are created. Clearly, such a configuration would be impossible for cubic symmetry, or more generally with only two independent direction of membranes, since overlapping corners in this case already imply an overlapping bulk.

The last step in showing that the honeycomb model considered in this chapter has extensive ground state degeneracy is to construct a state with an extensive number of such flippable local motifs. Such a state is shown in Fig. 4.3 (c). It is obtained from the Néel state by flipping lines of spins, costing zero energy as established in Fig. 4.2 (b). The lines flipped with respect to the Néel state in Fig. 4.3 (c) are indicated in gray. The bonds flipped in the local zero-energy-cost move are indicated in orange in the figure. They form a large hexagon, and flipping spins in such a pattern costs zero energy if the large hexagon is centered around any “flippable hexagon”, indicated in green. Note that while neighboring hexagons cannot be flipped simultaneously, the shown state is compatible with periodic boundaries and hence hosts an extensive number of independently flippable motifs. This implies a lower bound for the residual entropy

$$\lim_{T \rightarrow 0} S(T) := S_0 > \frac{N}{72} \log(2) \quad (4.10)$$

where $N = 2L^2$ is the number of sites. While the above construction shows explicitly that the number of ground states is exponentially large in the system size N , the resulting lower bound on the entropy S_0 is unlikely to be very tight. This is because the 24-spin move is not ergodic within the set of ground states, since there are many ways of combining membranes locally such that an even number of corners overlaps at any point.

4.3. Monte-Carlo Simulation of Fracton Phases

To study the behavior of the honeycomb model [Eq. (4.1)] also as a function of temperature, we develop in the following a novel cluster Monte-Carlo algorithm. Before explaining the algorithm in detail, we briefly discuss why conventional approaches fail.

4.3.1. Why local algorithms fail

The emergence of fractonic excitations in higher-order coulomb phases is fascinating, but it also presents a major complication when trying to study these phases numerically. Since fractons have restricted mobility, simple (local) Monte-Carlo algorithms will fail to equilibrate the system at low temperatures [15, 72, 75].

The central problem is that isolated fractons are immobile and instead one has to move pairs of excitations. Annealing out defects then requires at least four excitations to meet by moving pairs of defects around the system. Crucially, moving a pairs of defects generally requires a correlated flip of $\xi \sim e^{\beta J/3}$ spins, where ξ is the typical separations between defects (see previous section and Fig. 4.2). Analogously to the discussion of relaxation dynamics by quantum fluctuations in Ref. Sec. 2.3, this means that any local Monte-Carlo algorithm will have an autocorrelation time diverging rapidly at low temperature

$$\tau \sim \exp[\alpha e^{\beta J/3}] \quad (4.11)$$

for some constant $\alpha > 0$.

4.3.2. Moving fractons in pairs

As explained above, fractons have very restricted mobility. They can only move in pairs, and a given pair can only move in a fixed direction. Furthermore, moving a pair of fractons that is a distance d apart necessitates flipping $\mathcal{O}(d)$ spins and hence this constitutes a nonlocal update of the spin configuration. In the following, we present a cluster Monte-Carlo update rule fulfilling detailed balance which is able to move pairs of fractons by one step in their direction of mobility.

We define two kind of cluster moves, called the *parallel* and *perpendicular line update*. These are directly inspired by the creation of fractons by flipping membranes of spins with respect to a ground state as shown in Fig. 4.2 (b) and (c), respectively. The updates are performed as follows:

1. Chose a random bond $b_0 = \langle ij \rangle$. Initialize the cluster as $\mathcal{C}_0 = \{b_0\}$.
2. Attempt to flip the cluster \mathcal{C} with metropolis probability $p = \exp(-\beta\Delta E)$, where ΔE is the energy difference when flipping all spins in \mathcal{C} . If updated is accepted, go to 4.
3. If the update is rejected add another bond b' to the cluster:
 - a) If performing the parallel line update, this bond will be the next bond in the lattice reached by going in parallel to the last bond added to the cluster, see Fig. 4.2 (b).
 - b) If performing a perpendicular update, this bond will be either the next bond reached when going perpendicular to the cluster if the number of bonds in \mathcal{C} is odd, or the next-nearest bond in the same direction, when the number of bonds in \mathcal{C} is even. See Fig. 4.2 (c).

Then go back to 2.

4. Accept the full cluster update with acceptance probability

$$p_{\text{acc}} = \prod_{n=1}^{M-1} \frac{1 - \min[1, e^{-\beta(\Delta E_{M-n} - \Delta E_M)}]}{1 - e^{-\beta\Delta E_n}} \quad (4.12)$$

where ΔE_n is the energy difference in the n-th iteration of the loop above.

The above procedure implements exactly the movement of fractons as shown in Fig. 4.2 (b) and (c). The cluster \mathcal{C} constructed in step 3 corresponds to one row of a flipped membrane. At low temperature, such a strip will be flipped if either its endpoints coincide with the position of a configuration of fractons, or if it wraps around the system. The former case will move a pair of fractons, while the latter case changes ground nonlocally. In order to ensure that the construction of the cluster always terminates, we remove a bond from the cluster if it would be added a second time. Then, a cluster that wraps around the system twice is empty and hence $\Delta E = 0$.

The acceptance probability in Eq. (4.12) accounts for the probability of choosing a certain cluster and is needed to ensure detailed balance [123].

To benchmark our algorithm and as a consistency check of ergodicity, we show in Fig. 4.4, the internal energy of Eq. (4.1) as a function of Monte-Carlo time, when starting from a random initial state at a temperature of $T = J/5$. This temperature




Figure 4.4.: Evolution of the internal energy as a function of Monte-Carlo time, using the line-update at $T = J/5$, using the line update described in the main text. Dashed lines indicate the results from parallel tempering simulations (cf Fig. 4.6). Solid lines are an exponential fit to the late-time relaxation $E - E_{\text{eq}} = \exp(-t_{\text{MC}}/\tau)$. The relaxation time scale τ is shown as a function of system size in the inset.

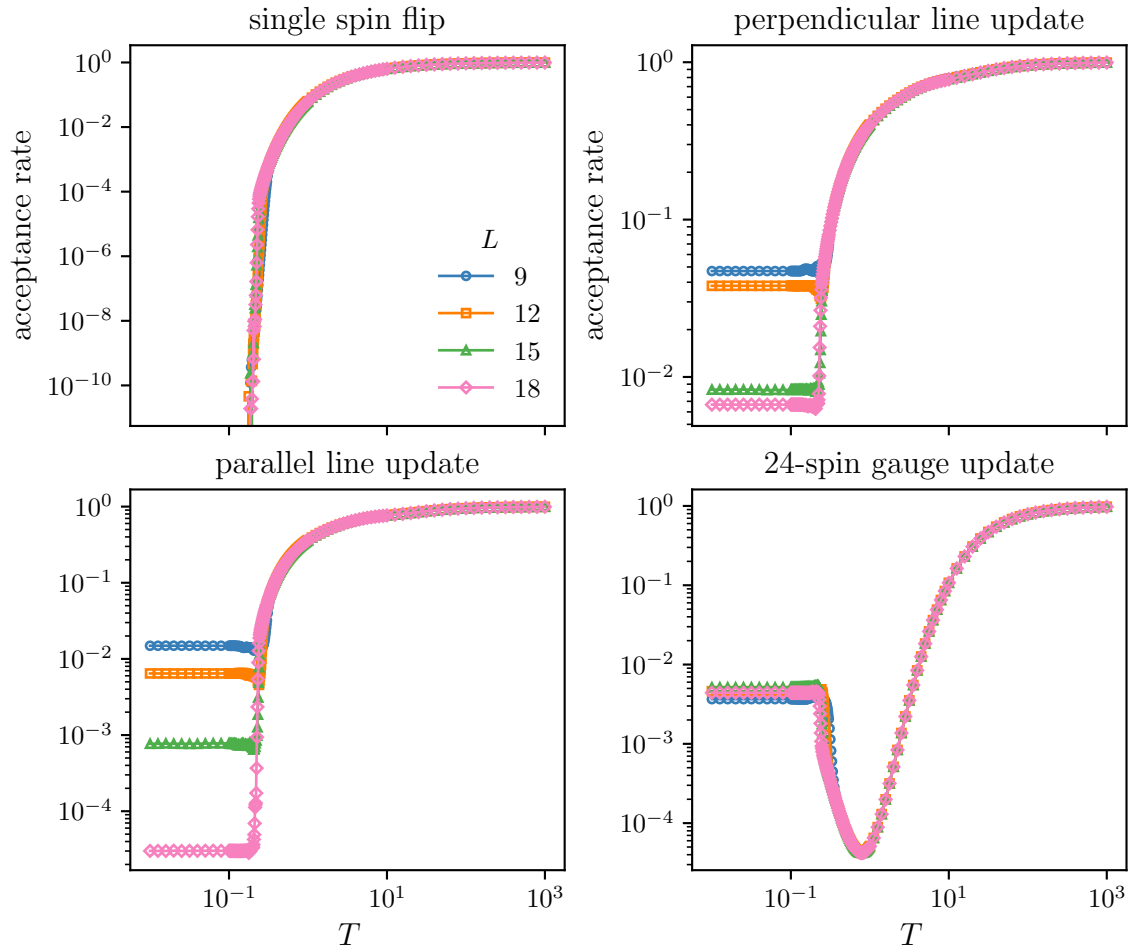


Figure 4.5.: Acceptance rate of the different moves in the honeycomb model as a function of temperature. The single-spin-flip move and the 24-spin gauge moves are local moves, while the perpendicular and parallel line moves are generally nonlocal cluster updates.

is well in the ground state regime, with much less than one excitation present on average. The data between times $t_{\text{MC}} = 10^m$ and $t_{\text{MC}} = 1.5 \times 10^{m+1}$ is obtained by alternating 9×10^m Monte-Carlo sweeps with 10^m measurements and repeating this 15 times. The data is averaged over 5000 runs for $L = 6, 9$ and over 2000 runs for $L = 12$. We compare the long time limit of this quench dynamics with the results from parallel tempering simulations (see Sec. 4.4 for full results of this). The long time limit of the energy after the quench agrees within error bars with the data obtained from parallel tempering. By fitting the long-time tail of the energy to an exponential form $E - E_{\text{eq}} \sim \exp(-t_{\text{MC}}/\tau)$, we extract a relaxation time τ that is shown as a function of system size L in the inset of Fig. 4.4. While the relaxation time still shows a significant dependence on system size, the cluster algorithm still poses a major improvement over any local version in practice. To demonstrate this, we performed a simulated annealing simulation using 100 intermediate temperatures between $T = J$ and $T = J/10$ and 10^6 sweeps per temperature. When using a spin-flip metropolis algorithm at $L = 18$, out of 500 runs not a single one manages to anneal out all defects. In contrast, our cluster algorithm in the same configuration reaches zero energy in all cases. This is because as shown in Fig. 4.5, the acceptance probability of single spin flips is essentially zero at low temperatures $T < 0.2J$. In contrast, the acceptance probability of the two cluster moves (the perpendicular and parallel line update) is nonzero. The acceptance rate of both moves still decreases as system size is increased, which is presumably the reason for the significant system-size dependence of the relaxation time observed in Fig. 4.4. We also implement the minimal local move, that is the 24-spin move shown in Fig. 4.3 (c). As expected for a local move, its acceptance rate is roughly independent of system size. Since it connects different ground states at zero energy cost it also has a nonzero acceptance rate even at the lowest energies.

4.4. Thermodynamics and Correlations

In order to reliably equilibrate large systems, we use the cluster algorithm discussed in the previous section, together with the local move shown in Fig. 4.3 (b-c), and augment this with feedback-optimized parallel tempering [124]. In Fig. 4.6, we show the internal energy, specific heat, as well as the entropy as a function of temperature for system sizes up to $L = 24$ (that is $N = 1152$ spins). To ensure equilibration, we compute the specific heat from both the fluctuation of the energy (markers) and as the derivative of the energy with respect to temperature (solid line in same figure) and check that the two estimates agree. This is considered “a stringent criterion for equilibration” and frequently used in simulations of model glass formers [125]. The entropy is obtained by integrating the specific heat from high temperature. For the residual entropy S_0 , we show the result of integrating both the variance of the internal energy (blue circles) as well as its derivative with respect to temperature (orange squares).

The data in Fig. 4.6 (a-c) shows clear signatures of a first-order transition from a high-temperature paramagnetic phase to a low-temperature phase with zero energy, that is $M_h = 0$ for all hexagons h . We estimate the thermodynamic residual entropy of this low-temperature phase by fitting the system size dependence of $S_0 = S(T \rightarrow 0)$ and extrapolate to $L \rightarrow \infty$. As shown in Fig. 4.6 (d), we use a linear fit to the entropy of the largest four system sizes considered as well as a quadratic fit to the

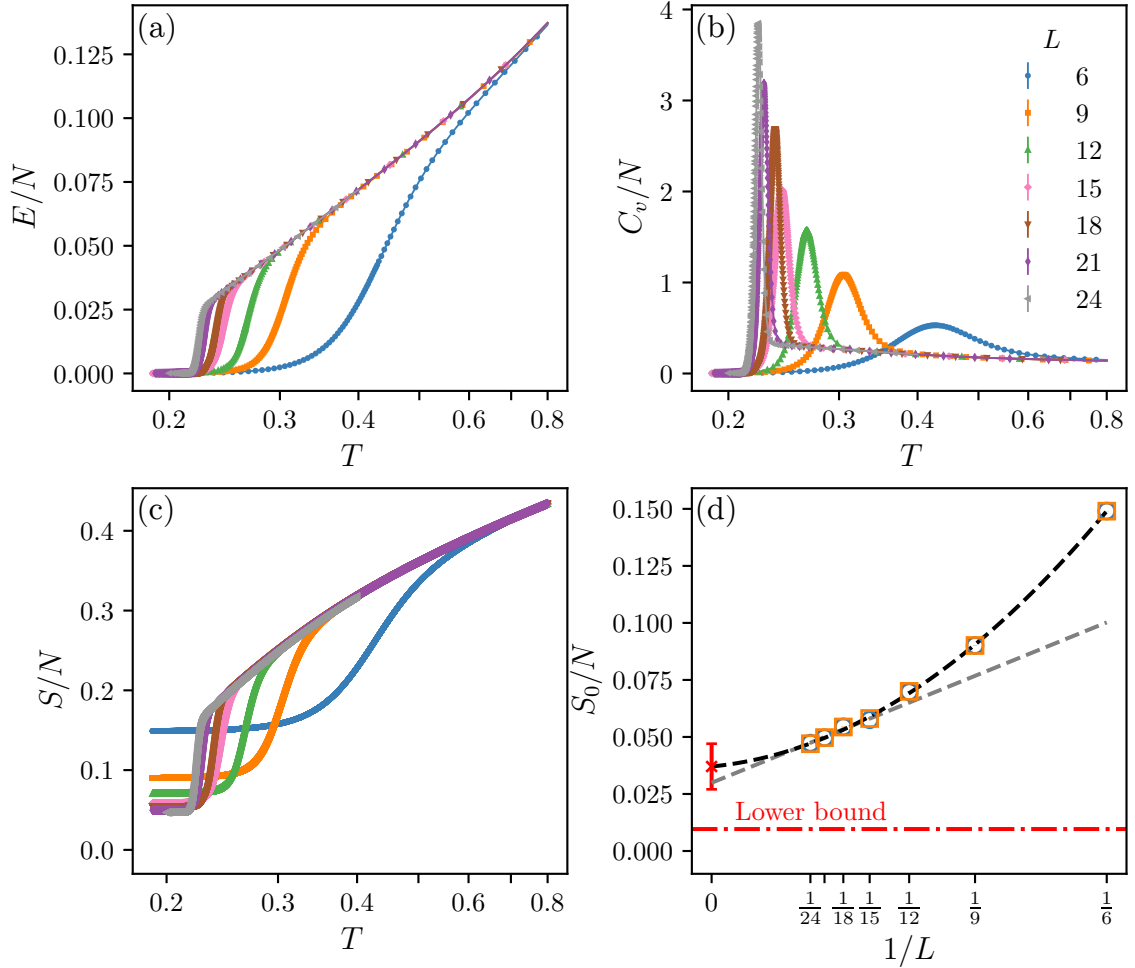


Figure 4.6.: Internal energy (a), specific heat (b), and entropy (c) of the honeycomb model [Eq. (4.1)] and a function of temperature, as obtained from Monte-Carlo simulations. For clarity, only small fraction of the simulated temperature range is shown, focussing on the apparent first-order transition. In (d), we show the residual entropy per site as a function of inverse linear system size. We also show linear (gray) and quadratic (black) extrapolations of the data to $L \rightarrow \infty$. The linear extrapolation is a fit to only the four largest system sizes.

full data. Taking the difference between the two as a measure of uncertainty, this yields an estimate of $S_0/N = 0.037 \pm 0.010$. This is quite a bit above the lower bound derived in Eq. (4.10), but comes with a large error bar due to significant finite size effects.

The structure factor within the low-temperature phase is shown in Fig. 4.7 and is fully consistent with a low-temperature higher-rank coulomb phase. It shows sharp four-fold pinch points at the zone boundaries for all system sizes considered. There are no sharp features such as Bragg peaks that scale with system size. This is evident from the line cut along a path in momentum space through high symmetry points, shown in Fig. 4.7 (b). The sharp peak present at the second zone center Γ' (corresponding to antiferromagnetic Néel order) for small system sizes $L < 18$ decreases in intensity rapidly as system size is increased. The structure factor of the Ising model also shares many qualitative features with the structure factor of the corresponding three-component vector spin model in Ref. [93], most prominently its strong anisotropy around the pinch points.

The presence of a sharp transition as a function of temperature is somewhat surprising since in two dimensions, the higher-rank coulomb liquid expected to describe the ground state regime is continuously connected to the paramagnet. It is however not inconsistent with a low-temperature liquid phase since a first-order transition is always possible also between continuously connected phases as demonstrated by the famous transition between the gaseous and liquid state of water. An alternative possibility is that the low-temperature phase is a so called “fragmented liquid”, that is the set of ground states, although extensive, breaks some symmetry on average. This is akin to the entropic order present in Onsager’s Hard rods [20] (cf also Sec. 2.1.1) and as discussed in Sec. 2.1.2, the ordered moment of such entropic ordering in magnets can be remarkably small [33]. A possible hint in this direction is that the maximum of the structure factor along the line cut shown in Fig. 4.7 scales roughly with linear system size. However, even though the system sizes studied are parametrically large they still show significant finite size effects as is evident from the system size dependence of the residual entropy. Ultimately, it is impossible to exclude the possibility of fragmentation in the absence of a more efficient algorithm and we leave this question open for future studies.

4.5. Conclusion and Outlook

In summary, this chapter constitutes a detailed study of what is to the best of our knowledge the first classical Ising fracton spin liquid. We explicitly showed that the model as defined in Eq. (4.1) has a number of ground states that grows exponentially with the number of sites and that its defects include fractons as well as lineons. Using a novel cluster Monte-Carlo algorithm that moves pairs of defects, we were able to study also the thermodynamic properties of the model as a function of temperature. Energy and specific heat are both consistent with a first-order transition from a high-temperature paramagnetic phase to a low-temperature phase which fulfills the ground state constraint of Eq. (4.1). We show that correlations in this low-temperature phase are consistent with a rank-2 coulomb liquid, showing four-fold pinch points but no other sharp features. An open question is the origin of the first order transition, and as a related question to what extent the liquid might be weakly fragmented.




Figure 4.7.: Left: Structure factor of the honeycomb model [Eq. (4.1)], computed for a $L = 24$ system at $T = J/5$. Right: structure factor around a line cut as function of system size.

The discovery of a relatively simple Ising model, with finite-range, two-body interactions establishes a useful platform for the further exploration of fractonic physics. This could include the perturbative introduction of quantum effects via transverse fields or transverse exchange. This may be a better setting in which to study quantum effects on fractons than in the Heisenberg models suggested in Ref. [93], for which numerical calculations suggest that quantum fluctuations wash out the multifold pinch points [126]. Here, the emergent Gauss's law is protected by a finite gap, so it may be more robust. Even if instanton effects drive the emergent gauge theory into a confined phase (as they do for the ordinary U(1) gauge theory in 2+1 D), the low temperature physics can still show interesting features related to the liquid phase.

Our purpose-built Monte Carlo algorithm provides a template for future numerical studies of Type-I fractonic models. Future work could address in more detail the topics of relaxation and disorder-free glassiness. The successful demonstration of an Ising fracton spin liquid, based on a Hamiltonian originally constructed for continuous spins [93], also raises the question of whether other classical spin liquids with higher-moment conservation laws [113, 114] have Ising realizations, and what their properties may be.

Given the rapid progress in physical simulation platforms based on Rydberg atoms and superconducting qubits, the dynamics of fractons in our model may not be far from exploration in the laboratory.

Part II.
Artificial Systems

5. Quantum Error Correction and Decoding

The idea of quantum computing has generated significant attention well beyond the physics community due to the possibility that it can solve specific problems more efficiently than classical computers [127–132]. Quantum computers are not the first alternative model of computation to be proposed. For example, Schönhage showed in 1976 that a classical computer able to process continuous variables can solve efficiently problems in the complexity class PSPACE [133]. This class consists of all problems whose memory requirement grows polynomially in the problem size, and it includes hard problems of practical interest such as the traveling salesman. The reason large-scale analog computers were never built is that there is no known fault-tolerant way of implementing them: errors accumulate during the computation and in contrast to the case of information stored using discrete variables, there is no known way of correcting these faults.

As Peter Shor demonstrated in 1995, the situation is different for quantum computation [134]. This is due to the fact that quantum mechanics is a juxtaposition of continuous variables and discrete observables. As explained in this chapter, it is the latter that enables the miracle of quantum error correction (QEC). For the purpose of this thesis, we use Shor’s nine-qubit code and Kitaev’s toric code as illustrative examples. A family of close cousins of the toric code, that is hyperbolic surface codes, will be the subject of the work presented in chapter 6.

5.1. From Classical to Quantum Error Correction

5.1.1. Classical repetition code

The simplest form of error correcting code is the one many of us use intuitively: to make ourselves understood in a noisy environment, we repeat what we said multiple times. This is formalized by the concept of the logical bit \bar{b} , which in the classical repetition code is defined as

$$\bar{b} = \underbrace{bbb\dots b}_{n \text{ times}} \quad \text{for } b \in \{0, 1\}. \quad (5.1)$$

Now, if the physical data is corrupted we can recover the original information by majority voting: if some of the bits b_j of \bar{b} are flipped, we count the number of 0s and 1s in \bar{b} and set all bits to whichever value appears more often. This procedure successfully restores the original state if less than half the bits have been flipped and fails otherwise. Assuming a bit flip rate $p < \frac{1}{2}$ for each physical bit and odd n the probability of failure is

$$p_L = F\left(\frac{n}{2}; n, 1-p\right) < \exp\left[-n D\left(\frac{1}{2} \middle| 1-p\right)\right] \quad (5.2)$$

where $F(k, n, p)$ is the cumulative distribution function of the Bernoulli distribution, i.e. the probability of having less than $n/2$ bits *not* flipped. In the last step, we used the Chernoff bound [135] and

$$D(a|b) = a \log\left(\frac{a}{b}\right) + \log\left(\frac{1-a}{1-b}\right) \quad (5.3)$$

is the Kullback-Leibler divergence between an a -coin and a b -coin. Most importantly, Eq. (5.2) shows that for $p < \frac{1}{2}$ the logical error rate p_L is exponentially suppressed in the code size n , while the overhead of implementation is $\mathcal{O}(n)$. For $p > \frac{1}{2}$, successful recovery is not possible and the local failure error rate p_L even grows as n is increased. The critical local failure rate $p_{\text{th}} = \frac{1}{2}$, below which error correction is possible, is called the code threshold and will be of central interest later on.

5.1.2. Shor's nine-qubit code

Implementing a procedure analogous to the one described in the previous section for preserving quantum information involves two major challenges. First, quantum information as mentioned before is continuous in nature, its fundamental unit being the state of a qubit

$$|\phi\rangle = \alpha |0\rangle + \beta |1\rangle \quad (5.4)$$

rather than the discrete bit b . Naturally, this implies a continuous set of possible errors, which complicates the analysis. Second, when correcting errors in the repetition code one measures the state of each constituent bit. For a qubit encoded in some larger set of degrees of freedom, measuring the state of all constituents generally would destroy their collective state. In the following, we illustrate how both of these issues can be solved using Shor's nine-qubit code [134], which corrects an arbitrary single-qubit error, as an example.

Quantum bit-flip code

A simple generalization of the classical repetition code encodes one qubit (using $n = 3$ for simplicity) as

$$|\bar{\phi}\rangle = \alpha |000\rangle + \beta |111\rangle. \quad (5.5)$$

The logical state $|\bar{\phi}\rangle$ lives in a subspace of the full Hilbert space spanned by the two basis states $|\bar{0}\rangle = |000\rangle$ and $|\bar{1}\rangle = |111\rangle$. This subspace is called the codespace \mathcal{C} . Instead of measuring the individual states of each qubit, we instead measure a set of commuting parities

$$S = \{Z_1 Z_2, Z_2 Z_3\} \quad (5.6)$$

Importantly, for any valid $|\phi\rangle \in \mathcal{C}$ measuring these operators yields $+1$ and preserves the state. The measurement outcome hence confirms that the state is part of the codespace, but without giving any additional information. This is crucial, since learning any information about the amplitudes α, β of the logical state $|\phi\rangle$ itself necessarily destroys it. Also, we can now correct any single bit-flip X_j , $j = 1, 2, 3$ since each of them anticommutes with (and hence changes the eigenvalue of) a unique combination of the parity measurements. To be explicit, we below summarize the outcome of the parity measurements $Z_1 Z_2$ and $Z_2 Z_3$, as measured in the state $X_j |\phi\rangle$ for $|\phi\rangle \in \mathcal{C}$:

Error	/	X_1	X_2	X_3
$Z_1 Z_2$	+1	-1	-1	+1
$Z_2 Z_3$	+1	+1	-1	-1

Assuming that at most a single bit flip occurred we thus learn its position, again without learning anything about the logical state $|\phi\rangle$ itself. The set of parity measurements that yields -1 is also called the syndrome of the error.

Correcting an arbitrary single-qubit error

Of course, bit flips are not the only possible error process, even on the single-qubit level. A general single-qubit error process can be modeled as

$$|\psi\rangle \rightarrow U|\psi\rangle = c_0|\psi\rangle + c_1X|\psi\rangle + c_2Z|\psi\rangle + c_3XZ|\psi\rangle \quad (5.7)$$

where c_i are complex constants. While this is a continuous set, performing measurement on the erroneous state discretizes the set of possible errors. To see this, consider the following encoding of one logical qubit into nine physical qubits

$$|\bar{\psi}\rangle = \alpha(|000\rangle + |111\rangle)^{\otimes 3} + \beta(|000\rangle - |111\rangle)^{\otimes 3}. \quad (5.8)$$

This is Shor's original nine-qubit code. It is the concatenation of the bit flip code discussed above with a so called phase-flip code

$$|\bar{\phi}\rangle = \alpha|+++ \rangle + \beta|--- \rangle \quad (5.9)$$

where $|\pm\rangle = |0\rangle \pm |1\rangle$. The phase flip code is implemented by measuring X -parities X_1X_2 , X_2X_3 and, analogously to the bit flip code discussed above [Eq. (5.5)], can correct a single phase flip Z_j . The full nine-qubit code is implemented by measuring the set

$$S = \{Z_1Z_2, Z_2Z_3, Z_3Z_4, Z_4Z_5, Z_5Z_6, Z_6Z_7, Z_7Z_8, Z_8Z_9, \\ X_1X_2X_3X_4X_5X_6, X_4X_5X_6X_7X_8X_9\}. \quad (5.10)$$

As is easy to check, all these operators mutually commute and their common $+1$ Eigenspace is given by the set of states $|\bar{\psi}\rangle$ in Eq. (5.8). The set of these operators is called the stabilizers of the code, and a large class of codes is defined by defining a set of stabilizers¹. The central insight is now that applying an arbitrary single-qubit unitary U on any qubit and measuring the stabilizers collapses the state into one of the three possibilities $X|\bar{\psi}\rangle$, $Z|\bar{\psi}\rangle$, or $XZ|\bar{\psi}\rangle$. This yields the discretization of errors promised above. Since as before any combination of single bit flip X_j and phase flip Z_j anticommutes with a unique set of stabilizers, the measurement outcome can be used to correct the now discretized error.

To summarize, Shor's code uses nine physical qubits ($n = 9$) to encode a single qubit, which is called the logical qubit. The logical qubit is more robust in the sense that we can correct an arbitrary single qubit error. This gain in robustness is usually quantified by the so called "distance" of the code. If a code can correct up to l single-qubit errors, we say it has distance $d = 2l + 1$. The three key properties of

¹Technically, the set S defines only the *generators* of the set of stabilizers but we will abuse wording here slightly in order to be less verbose.

a code, that is the number of physical qubits n , the number of logical qubits k and the distance d are often summarized as a triplet $[n, k, d]$. In this notation, Shor's code is a $[9, 1, 3]$ code.

While Shor's Code demonstrates the principle possibility of QEC, it is not very practical to implement because of the high-weight X stabilizers. The weight of these stabilizers even grows when considering generalization of the nine-qubit code to larger number of qubits n . Instead, current efforts focus on code families with local stabilizers such as surface codes, which I will introduce in the following.

5.2. Kitaev's Toric Code

The toric code is one member of a larger class of codes called surface codes, which can be defined on any tessellation of any orientable two-dimensional surface [14, 136, 137]. For concreteness, we focus in the following on the specific example of the toric code, but comment at some places on the more general case. This will be useful in the discussion of hyperbolic surface codes in chapter 6. This family of codes is constructed analogously to the toric code but have quite different properties because of their underlying hyperbolic geometry.

5.2.1. Stabilizers and codespace

The toric code is defined by putting qubits on the edges of a square grid and defining two kind of stabilizers defined on the vertices v and faces p of the lattice

$$A_v = \prod_{e \in \delta v} Z_e, \quad B_p = \prod_{e \in \partial p} X_e, \quad (5.11)$$

The first kind A_v is as a product of Pauli- Z operators over all edges incident on the vertex v . This set, denoted as δv , is also called the ‘‘coboundary’’ of the vertex v . The second kind, B_p , is defined as the product of Pauli- X operators over all edges in the boundary of the face p . This is illustrated in Fig. 5.1 (a). It is easy to check that all operators A_v, B_p commute. First, they commute trivially among themselves. Then, the coboundary of any vertex, δv , and the boundary of any face, ∂p , share either no edge or exactly two edges. This implies that the support of any A_v shares an even number of qubits with the support of any B_p . Since X_e and Z_e on the same edge e anticommute, this already implies that all A_v commute with all B_p .

As in Shor's nine-qubit code, we now define the codespace \mathcal{C} as the common +1 Eigenspace of all stabilizers $\{A_v, B_p\}$. As we will see below, if the code is defined on a grid with periodic boundary conditions, that is on a square tessellation of the torus, the resulting stabilizer code encodes two logical qubits $k = 2$.

As Kitaev already discussed in his original paper, the construction above can be generalized to any tessellation of an orientable, two-dimensional manifold M . Stabilizers in this case are defined exactly as above on vertices and faces and they commute by the same arguments. Remarkably, the properties of the codespace depend only on the topology of the manifold. In particular, the number of logical qubits is given by the genus g of M as

$$k = 2g(M) \quad (5.12)$$

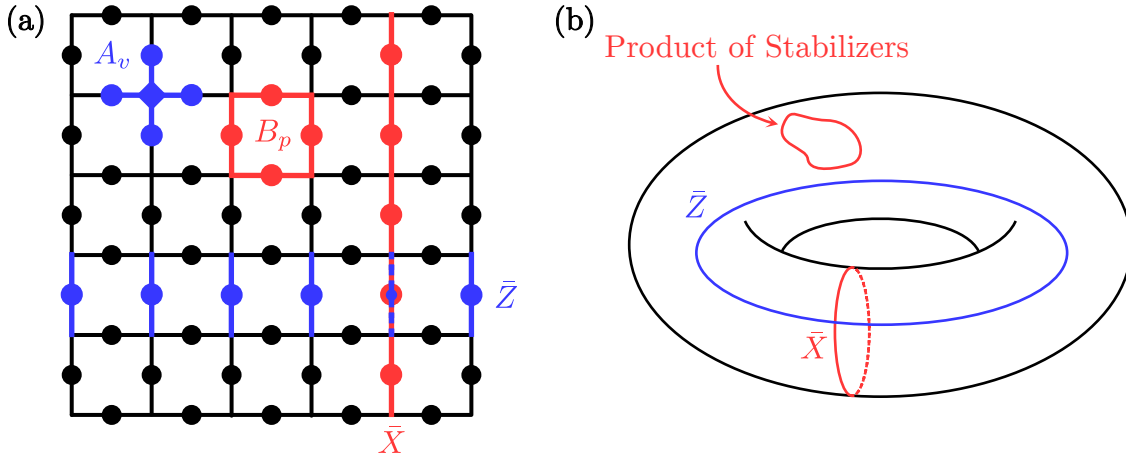


Figure 5.1.: Kitaev's toric code, with qubits on the edges of a square grid. It has two kinds of stabilizers A_s on vertices and B_p on faces. The logical Pauli operators \bar{X} and \bar{Z} are given as nontrivial loops wrapping around the periodic boundary conditions. Panel (a) shows the lattice construction while (b) illustrates the homological picture on the torus.

This relation follows quite naturally when constructing the logical Pauli operators of the code, which we do in the following.

5.2.2. Logical operators

The logical Pauli operators \bar{X}_j and \bar{Z}_j , $1 \dots k$ of the code are defined as a set of operators that act as Pauli operators on the logical subspace. By this we mean that first, they preserve the code space and second, they fulfill the Pauli algebra $[\bar{X}_i, \bar{X}_j] = [\bar{Z}_i, \bar{Z}_j] = 0$, and $\{\bar{X}, \bar{Z}\} = 0$ [138]. For stabilizer codes, they are given by products of X and Z operators respectively, that commute with all stabilizers but are not a product of any subset of them.

For the toric code, the logical operators \bar{X}_j (\bar{Z}_j) are given by non-contractable closed loops on the (dual) lattice. This is shown in Fig. 5.1 (a) on a 5×5 grid with periodic boundary conditions and illustrated on the torus in panel (b) of the same figure.

To see this, consider first the logical \bar{X} . By the same argument that we used to show that all stabilizers commute, any product of X operators along a closed loop on the lattice commutes with all stabilizers. A “closed loop” here is defined as any set of edges such that an even number of those edges is incident at any vertex. Products of the stabilizers B_p naturally are also products of X operators over closed loops. However, note that any product of B_p is given exactly as the product over the boundary of a set of faces

$$\prod_{p \in A} B_p = \prod_{e \in \partial A} X_e. \quad (5.13)$$

Hence, the logical operators \bar{X} are given by those closed loops which are not boundaries. These are precisely the non-contractible loops on the lattice. Note that since stabilizers act as the identity on any element of the codespace, logical operators are uniquely defined only up to multiplication with stabilizers. This is equivalent to say




Figure 5.2.: Decoding a bitflip error in the toric code. The error anticommutes exactly with those stabilizers located at the endpoints of the error chain E . Upon measurement, these will return -1 and form the syndrome. To correct the error, a recovery R is applied that pairs up the syndrome. Recovery succeeds if the combined action of recovery and error is a product of stabilizers, and fails if it is a logical operator.

that we can define \bar{X}_j as the product of X operators on any loop which wraps an odd number of times around the j th handle of the torus (and an even number of times around all other handles).

The derivation of the logical \bar{Z}_j operators follows completely analogously to the discussion above, but on the dual lattice. This can be seen by realizing that the edges in the coboundary δv^* of a dual lattice site v^* are given exactly by the boundary operator of the face p dual to v^* .

On the torus, there are exactly two distinct non-contractible loops and correspondingly two logical \bar{X}_j and two \bar{Z}_j operators which are not equivalent up to multiplication of stabilizers. Accordingly, the code has two logical qubits. A particular choice for one of the two sets is drawn in Fig. 5.1 (a).

The derivation above also helps to illustrate the general relation in Eq. (5.12) between the number of logical qubits k and the genus of the manifold the code is defined on. As discussed above, logical operators correspond exactly to those loops that are closed, but which are not boundaries of any set of faces. Logical operators are then equivalent up to multiplication with stabilizers. The study of the corresponding equivalence classes of closed loops is called homology and for this reason the toric codes and similar constructions are often called “homological codes”. For the purpose of this thesis, while we will borrow some nomenclature, we do not want to venture too deep into formal math. For our purposes, it suffices to know one central result of homology, namely that the number of non-nonequivalent closed loops on a tessellation of any two-dimensional, orientable manifold is exactly $2g$, where g is the genus of the manifold. Realizing that the number of logical qubits k is equal to the number of logical operator pairs and hence the number of nonequivalent loops, Eq. (5.12) follows immediately.

5.2.3. Decoding the toric code

Successful recovery from errors

When performing error correction in the toric code (and any other code for that matter), the idea is to measure all stabilizers periodically in time. Assuming that we started with our system being in a logical state $|\phi\rangle \in \mathcal{C}$, we can use the outcome of these periodic measurements to diagnose and correct local errors.

As discussed before, it is sufficient to correct local bit- and phase-flips, that is Pauli- X and Z errors. Focussing first on bit-flips X , consider as an example the chain (that is, for our purposes, a set of edges) of errors E shown in red in Fig. 5.2. It anticommutes exactly with those A_v stabilizers located at the endpoints (the boundary) of the error chain, i.e. those vertices where an odd number of edges of the chain are incident. These stabilizers upon measurement evaluate to -1 and are called the error syndrome, denoted as $S = \partial E$. Note that any local error which has no endpoints, such as the one shown on the top right of Fig. 5.2, does not produce a syndrome but is a product of stabilizers, so it acts as the identity on the logical subspace.

Given a syndrome S , the goal is now to apply a recovery operation R that takes the system back to the code space \mathcal{C} . Ideally, this is done with a high probability of returning to the exact same initial state $|\phi\rangle$. The procedure of finding such a recovery given a syndrome is called decoding. We can guarantee returning to the code space by using a recovery that pairs up syndromes, that is we choose the recovery such that its endpoints coincide with the syndrome, $\partial R = S$. In this case, the combination of error and recovery chains is a closed loop, $\partial(R \circ E) = \emptyset$, and it commutes with all stabilizers. Such a recovery operation is successful however, if and only if the combined action of error and recovery is a product of stabilizers, since only then it acts as the logical identity, $R \circ E = \bar{I}$. For the error in Fig. 5.2, such a recovery operation is shown in green. Recovery fails, if the combined action of error and recovery is a logical operator, as shown in Fig. 5.2 in orange.

Note that although we have here only discussed correcting X errors, when assuming independent X and Z errors we can correct both types of errors separately. By the same argument used for the construction of logical operators, the correction procedure for phase-flips Z is analogous to the above discussion but on the dual lattice.

Minimum-weight perfect matching

To summarize, a working decoder for the toric code takes as input a syndrome S , that is the list of stabilizer measurements that yielded -1 , and then proposes a recovery operation R such that first $\partial R = S$ and second such that with high probability $\partial R \circ E = \bar{I}$ is a product of stabilizers. In terms of the set of edges forming R and E , we say that they form a loop of trivial homology.

A simple heuristic to achieve this task is given by so called minimum-weight perfect matching (MWPM) [136, 139, 140]. The idea is to just implement the *shortest* recovery operation that pairs all syndromes. Assuming that phase- and bit-flip errors occur independently on each qubit with uniform probability p , this corresponds to choosing as a recovery the most likely error to have caused the given syndrome. Finding the shortest path pairing up a given set of vertices on a graph




Figure 5.3.: The importance of entropy. I shown an error on a $d = 6$ toric code, and two possible recovery operations. The operation indicated in green would be successful while the orange operation fails. While both possibilities have individually the same size, there are five equivalent choices (up to stabilizers) with the same size for the green recovery while the orange recovery is unique unless considering larger sizes. An optimal decoder would hence chose the green recovery.

constitutes a global optimization problem. Remarkably, in a paper with the lovely title “Paths, Trees, and Flowers” [139], Edmonds showed in 1965 that it is possible to solve this problem efficiently on any graph. The run time of the algorithm scales as $\mathcal{O}(n^3)$ in general but can be reduced to $\mathcal{O}(n)$ in the limit of small error rate [140].

For a toric code defined on a $L \times L$ grid, MWPM is naturally guaranteed to succeed if less than $L/2$ errors occurred between syndrome measurements. This is because in this case the maximum distance between syndromes is less than $L/2$ and the combination error and the minimum-weight matching is guaranteed to be a trivial loop. The distance of the toric code is the minimal size of a logical error and hence $d = L$. In the notation introduced at the beginning of the chapter, the toric code is hence a $[n, k, d] = [2L^2, 2, L]$ code.

A natural question to ask is what is the threshold of the toric code using MWPM. The threshold is the critical physical error probability p_{th} such that for $p < p_{\text{th}}$ the logical failure rate is suppressed exponentially in L . In contrast to the situation in the repetition code, where $p_{\text{th}} = \frac{1}{2}$, computing this value for the toric code is more complicated and either requires simulation of the code, or a mapping of the decoding procedure to a statistical mechanics model [136]. This mapping is quite insightful and will be presented in the next section. Before going there, it is however worth discussing the limitations of MWPM.

Beyond MWPM

MWPM is suboptimal in the sense that errors exist which are correctible in principle, but which are not corrected successfully by MWPM. This is because as already mentioned above, MWPM corresponds to choosing as recovery the most likely error to have caused this syndrome.

To see this, first note that just as any other operation on the code, also errors and recoveries are equivalent up to multiplication with stabilizers. An optimal decoder now chooses as a recovery not the most likely error E compatible with the syndrome, but instead chooses a recovery from the most likely equivalence class compatible with the syndrome. This is equivalent to maximizing the probability

that the combination of error and recovery is a trivial loop. Two errors E and E' are in the same equivalence class if they are related by some product of stabilizers s , that is $E = E' \circ s$. All errors in the same class have the same syndrome. For the toric code, and assuming we decode bit- and phase-flip independently, there are four such equivalence classes for each type of error, with representatives $E = I, \bar{P}_1, \bar{P}_2, \bar{P}_1 \bar{P}_2$, $P = X, Z$.

To see that optimizing over equivalence classes rather than single errors sometimes leads to different results, consider the case shown in Fig. 5.3. The figure shows an error on a slice of an $L = 6$ toric code, together with two different recovery operations in green and orange. Both have the same size so the MWPM decoder would choose one of them at random. However only one of them (green) corresponds to a successful recovery. While both have the same size and hence the same probability, their equivalence classes have different probabilities. In particular, there are five possible recoveries of the same size that are equivalent to the representative in green, while the orange recovery is unique unless considering larger weights. A optimal decoder thus chooses the green recovery deterministically, outperforming MWPM.

The case shown in Fig. 5.3 is arguably quite artificial and one can ask how much is gained by considering corrections to MWPM or even an optimal decoder. Assuming independent bit- and phase-flip errors, this question can be answered within the statistical mechanics mapping discussed in the next section. The answer will be that MWPM for the toric code under independent bit- and phase-flip noise is indeed almost optimal.

5.3. The Statistical Mechanics of Decoding

5.3.1. Optimal decoding as a statistical mechanics problem

Maximum likelihood decoding

Let us formalize the task of an optimal decoder by writing the probability of equivalence classes or errors causing a given syndrome. Considering only independent bit-flip noise with error rate p on n qubits, the probability of a specific class of errors to cause a syndrome S is given by

$$\bar{p}(E_0|S) = \sum_{s \in \langle \{B_p\} \rangle} p(s \circ E_0) = \sum_{s \in \langle \{B_p\} \rangle} p^{|s \circ E_0|} (1-p)^{n-|s \circ E_0|} \quad (5.14)$$

where E_0 is some representative of the class with $\partial E_0 = S$ and the sum is over all possible products s of stabilizers B_p . Finally, the probability of a specific error configuration $O = s \circ E_0$ is given by $p(O) = p^{|O|} (1-p)^{n-|O|}$, where the absolute value $|O|$ denotes the size of the support of the operator O , that is the number of qubits on which it does not act as the identity.

An optimal decoder, also called the maximum-likelihood decoder, has to compute the above probabilities \bar{p} for each equivalence class and chooses a recovery in the equivalence class that maximizes it

$$R_{\text{ML}}(S) = E_0 \circ \bar{L}_{\text{ML}}, \quad (5.15a)$$

$$\bar{L}_{\text{ML}} := \arg \max_{\bar{L}} [\bar{p}(E_0 \circ \bar{L}|S)], \quad (5.15b)$$

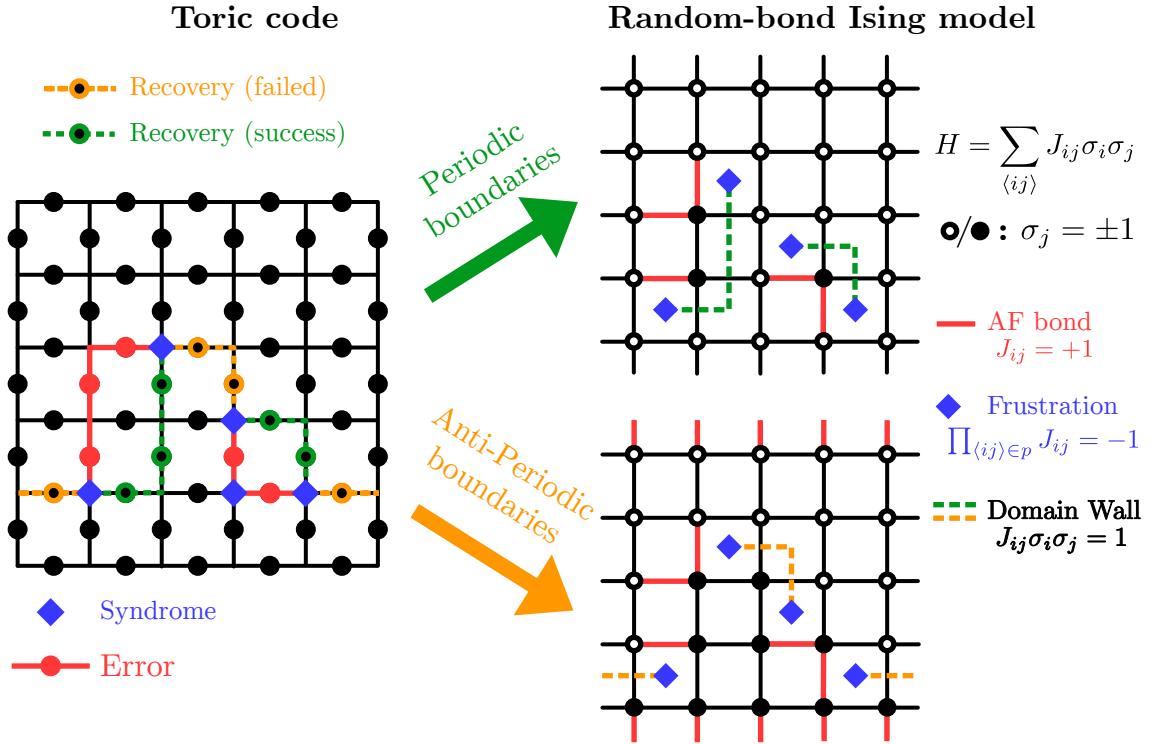


Figure 5.4.: Mapping between decoding of the toric code and the random bond Ising model (RBIM). Decoding the error shown in the left, the corresponding RBIM is defined on the dual lattice (right). Qubit errors correspond to antiferromagnetic bonds, the syndrome corresponds to frustrated plaquettes and possible recovery operations correspond to domain walls. Importantly, different equivalence classes of recoveries correspond to different boundary conditions in the RBIM. For the two examples shown in here in green and orange, these are periodic and anti-periodic boundary conditions, respectively.

where the \bar{L} denote all possible 2^k combinations of logical operators, including the identity $\bar{L} = \bar{I}$. Note that in the above expression, the maximum-likelihood recovery R_{ML} is (up to stabilizers) independent of the choice of E_0 , as long as $\partial E_0 = S$. First, multiplying any set of stabilizers to E_0 only changes R_{ML} up to stabilizers. Second, even choosing $E'_0 = E_0 \circ \bar{L}$ will result in $\bar{L}'_{\text{ML}} = \bar{L}_{\text{ML}} \circ \bar{L}$ and hence $R'_{\text{ML}} = R_{\text{ML}}$. In the following discussion we therefore just set $E_0 = E$, without loss of generality.

Statistical mechanics mapping I: partition function

As noted already before, the set of all possible products of stabilizers, $\langle\{B_p\}\rangle$, is in one-to-one correspondence to the set of all boundaries on the lattice, $\text{Im } \partial$). The central idea of the statistical mechanics mapping is now to write what is effectively a sum over boundaries in Eq. (5.14) as a sum over domain walls in a suitably defined Ising model.

As is illustrated in Fig. 5.4, the corresponding Ising model is defined on the dual lattice, such that each site j in the Ising model corresponds to a face $p(j)$ in the original lattice and each edge $\langle ij \rangle$ in the dual lattice corresponds to an edge $e(i, j)$ in the original lattice. Any product of stabilizers s is then in one-to-one correspondence with one configuration of Ising variables $\{\sigma_j\}$. The error and its syndrome are naturally encoded in the Ising couplings J_{ij} , as will become clear in the following. Specifically, for $s = \prod_{p \in A} B_p$ we define $\sigma_j = +1$ if $p(j) \in A$ and $\sigma_j = -1$ otherwise, and also $J_{ij} = +J$ if $e(i, j) \in E$ and $J_{ij} = -J$ otherwise. Then, the probability \bar{p} in Eq. (5.14) can be rewritten as the partition function of an Ising model

$$\bar{p}(E|S) \propto Z_{\text{RBIM}} = \sum_{\{\sigma_j\}} e^{-\beta H_{\text{RBIM}}[\{J_{ij}\}]}, \quad (5.16)$$

$$H_{\text{RBIM}}[\{J_{ij}\}] = \sum_{\langle ij \rangle} J_{ij} \sigma_i \sigma_j. \quad (5.17)$$

$$\beta = \log\left(\frac{p}{1-p}\right), \quad (5.18)$$

Since the proportionality constant in Eq. (5.16) is independent of the choice of E , we can reinterpret the task of the optimal decoder to compute the partition function (or equivalently, the free energy) of an Ising model on the dual lattice, with a specific configuration of couplings $\{J_{ij}\}$, $J_{ij} = \pm 1$. The Hamiltonian in Eq. (5.17) is also known as the random bond Ising model (RBIM), and was first introduced by Edwards and Anderson to model the interaction of dilute magnetic alloys [41]. The condition on the inverse temperature $\beta = T^{-1}$ in Eq. (5.18) in this context is called the Nishimori condition [141] and defines a line in the two-dimensional $p - T$ phase diagram of the model. This Nishimori line has been studied quite intensely because of an emergent additional symmetry which allows for derivation of some exact results (see Ref. [141] for details).

Statistical mechanics mapping II: The full dictionary

To understand this mapping in more detail, consider a specific example as shown in Fig. 5.4. There, I show on the left a 5×5 patch of the toric code with an error E indicated in red, the corresponding syndrome indicated by blue diamonds, and

two recovery operations in green and in orange. The two recovery operations are from different equivalence classes, that is they differ by a nontrivial loop wrapping horizontally around the torus. Consider now the mapping as outlined above, first for $E_0 = E$. The corresponding random-bond Ising model is shown on the top right of the same figure. Antiferromagnetic bonds $J_{ij} = 1$ are shown in red, and blue diamonds mark those plaquettes p corresponding to a syndrome in the original lattice. Given the disorder configuration these turn out to be exactly the “frustrated plaquettes”, defined as

$$\prod_{\langle ij \rangle \in p} J_{ij} = -1 \quad (5.19)$$

which means that no configuration of Ising spins around the plaquette can minimize all bond energies simultaneously. As Toulouse showed 1977 [23], disorder configurations in the RBIM are thermodynamically equivalent up to their configuration of frustration. This corresponds naturally to the fact that errors within the same equivalence class are equivalent as long as they cause the same syndrome.

The configuration of the Ising variables indicated in Fig. 5.4 by full and empty circles is the ground state of the model. The ground state, even in the presence of bond disorder, can be computed efficiently on planar graphs [142–144] by performing minimum-weight perfect matching of frustrated plaquettes. Indeed, the domain walls in the ground state, that is the set of bonds with energy $J_{ij}\sigma_i\sigma_j = +1$, correspond exactly to the recovery shown in green. The green recovery however is *not* the minimum-weight perfect matching of the syndromes. This would be the other recovery shown in orange on the left. The reason for this discrepancy is, of course, that the torus is not a planar graph. The ground state of the model in this case corresponds not to the overall minimum-weight matching but to the matching within a specific equivalence class of matchings, given by the boundary conditions of the model.

To see this, consider the same Ising model, but with anti-periodic boundary conditions around the vertical handle, shown in the bottom right of Fig. 5.4. The ground state is again indicated by filled and empty circles, and in this case the domain walls are given indeed by exactly the matching shown in orange. The choice of antiperiodic boundary conditions of course corresponds exactly to choosing as the trial error not E , but instead $E \circ \bar{X}_1$ where \bar{X}_1 is the logical X -operators wrapping horizontally around the torus.

In summary, we can actually reinterpret both, the minimum-weight perfect matching decoder as well as the maximum-likelihood decoder in terms of the statistical mechanics mappings in Eq. (5.16). Minimum-weight perfect matching proposes a recovery corresponding to the most likely error. In the statistical mechanics model, this corresponds to minimizing the energy over different boundary conditions. Likewise, the maximum likelihood decoder chooses a recovery from the most likely equivalence class of errors. On the statistical mechanics side of the mapping, this corresponds to minimizing the free energy over all boundary conditions.

Note that using this formulation of the mapping, any method of approximating the (free) energy of the RBIM for a given disorder realization readily yields a decoder for the toric code as follows. Given some syndrome, choose some matching of the syndrome. This matching does not have to be optimal and hence it can be obtained more efficiently than by running MWPM. This matching will serve as the trial error E_0 , which translates to the disorder realization J_{ij} . Now, minimizing the

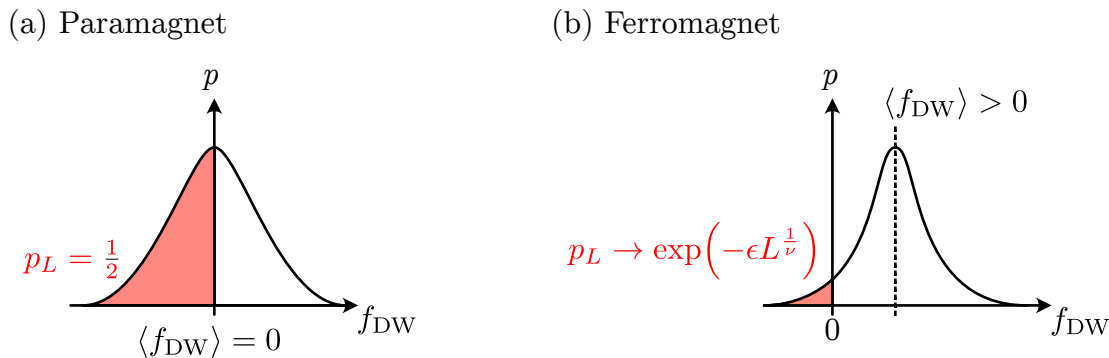


Figure 5.5.: Threshold and phase transition. Decoding being successful corresponds to periodic boundary conditions having the lowest free energy. This corresponds to a positive domain wall free energy f_{DW} . In the paramagnet (a), its expectation value is zero, $\langle f_{\text{DW}} \rangle = 0$, and decoding fails with a rate $p_L = 1/2$ as $L \rightarrow \infty$. In contrast, in the ferromagnet (b) the expectation value is nonzero $\langle f_{\text{DW}} \rangle > 0$ and its distribution becomes more narrow as $L \rightarrow \infty$, suppressing p_L exponentially.

approximate free energy over all possible 2^k boundary conditions yields a guess for R_{ML} . As the approximation is made more and more accurate, such a decoder will become more and more optimal.

5.3.2. Thresholds and phase transitions

A particularly interesting property of codes, but importantly also of decoders, is their threshold. As mentioned before, this is the maximal error rate p_{th} for which successful decoding is possible. More importantly, for $p < p_{\text{th}}$ the logical error rate p_L for which decoding fails is exponentially suppressed as the code size n is increased. Generally we expect

$$p_L \rightarrow \exp(-\epsilon d^\phi) \quad \text{for } p < p_{\text{th}} \text{ and } d \rightarrow \infty \quad (5.20)$$

for some constants $\epsilon, \phi > 0$ and d the distance of the code.

As we show in the following, for the toric code with distance $d = L$, the threshold is given by the critical value of disorder of the random bond Ising model [Eq. (5.17)]. For the maximum likelihood decoder, the relevant critical point will be that along the Nishimori line [Eq. (5.18)], while the threshold for minimum-weight perfect matching is given by the critical value of disorder at zero temperature. The value of ϕ in Eq. (5.20) in both cases will be given by $\phi = 1/\nu$, where ν is the critical exponent of the correlation length.

When decoding can't succeed

To derive the results stated above, one has to understand why it becomes impossible for a decoder to succeed on average at large error rate p more quantitatively. For this, it is helpful to recall the explicit construction of the ML decoder in Eq. (5.15). The optimal decoder chooses as a recovery the most probable equivalence class or errors. Taking without loss of generality the trial error to be the real error $E = E_0$ (see previous section), the threshold condition in Eq. (5.20) hence is equivalent to

the most probable equivalence class being the trivial class $l = 0$, with probability approaching one exponentially.

On the statistical mechanics side of things, this corresponds to the free energy being minimized by periodic boundary conditions, while the free energy of all other choices increases at least linear with the system size L . Quite remarkably, this corresponds exactly to the system being in the ferromagnetic phase. More specifically, consider the domain-wall free energy of the RBIM

$$f_{\text{DW}} = \frac{(F_{+-} - F_{++})}{L} \quad (5.21)$$

where F_{++} is the free energy of the RBIM with periodic boundary conditions and F_{+-} is the same free energy but for anti-periodic boundary conditions around one of the handles. Averaging over disorder realizations $\{J_{ij}\}$, the random-bond Ising model is known to undergo a phase transition along the Nishimori line [Eq. (5.18)], from a large-disorder (and high-temperature) paramagnetic phase to a low-disorder (and low temperature) ferromagnetic phase. The domain wall free energy f_{DW} serves as an order parameter to distinguish the two phases, with $\langle f_{\text{DW}} \rangle = 0$ in the paramagnet and $\langle f_{\text{DW}} \rangle > 0$ in the ferromagnet.

The probability of decoding to fail is now exactly the probability to draw an error corresponding to a disorder realization with negative domain wall free energy. The distribution of f_{DW} for different disorder realization is sketched in Fig. 5.5. In the paramagnet, its expectation value is expected to be symmetric around zero for large L , leading to $p_L = 0$ as shown in Fig. 5.5 (a). In the ferromagnet, sketched in panel (b) of the same figure, the distribution is instead centered around a finite expectation value, with tails reaching also to negative values at finite L . However, assuming self-averaging for the moment, the distribution becomes narrow as $L \rightarrow \infty$, and the weight of its tails is exponentially suppressed.

The discussion here is for a specific non-trivial class, that is for one of the possible logical errors. However, since the choice in Eq. (5.21) is arbitrary, we expect that after averaging over disorder, all possible logical errors are exponentially suppressed in system size. Since the number of possible logical errors, $2^k - 1$, is constant, the maximum-likelihood threshold of the toric code is then given by the position of the critical point of the RBIM along the Nishimori line.

The assumption of self-averaging in a magnet with large disorder, originally invented as a model for spin glasses (!)[41], may seem quite egregious. However, it can be shown that there is no spin glass phase along the Nishimori line [141] and self averaging is expected in this case. It is noteworthy however that, even if the long-distance behavior is rather conventional, there can be strong finite size effects [145, 146]. These are naturally relevant for codes, where one is typically interested in large, but not *extremely* large instances, and certainly not in codes using a number of qubits comparable to the number of atoms in typical samples of real materials.

In the case of the toric code and the RBIM specifically, the critical point on the Nishimori line is quite well-studied, originally because of its relation to quantum Hall physics [147, 148]. It is called the Nishimori point and its position and critical exponents are well established numerically [148]

$$p_c = 0.1093 \pm 0.0002, \quad (5.22)$$

$$\nu = 1.50 \pm 0.03. \quad (5.23)$$

From general scaling arguments, one then expects for $L \rightarrow \infty$

$$p_L = \exp\left(-\alpha(p_c - p)L^{\frac{1}{\nu}}\right), \quad (5.24)$$

for some constant α .

Finally, while the discussion above focused on the optimal decoder, rather than minimum-weight perfect matching, the discussion naturally generalizes to this case. To this end, we simply take the zero temperature limit in which case the free energy becomes equal to the energy. Hence, the threshold of MWPM is given by the zero temperature phase transition of the RBIM. From the code perspective it is not surprising that the numerical value of this threshold is smaller than that of the optimal decoder $p_{c,0} < p_c$. It is less natural when considering the statistical mechanics model, since it implies that the phase boundary curves *inward* as the temperature is lowered below the Nishimori point. This unusual shape of the phase boundary however was a known feature of the model called “reentrance” and can be established by exact analytical arguments [141]. The most precise estimate of $p_{c,0} = 0.1031 \pm 0.001$ comes from simulating the decoding of the toric code [149]. Comparison with Eq. (5.23) shows that MWPM matching for the toric code, although not optimal in principle, is almost optimal at least in terms of its threshold value.

5.3.3. Beyond the toric code and phenomenological noise

While the discussion of the statistical mechanics mapping above took the decoding of the toric code under phenomenological noise (that is independent bit- and phase-flip noise) as an illustrative example, similar mappings can be developed for both other codes [150–152] and different noise models [153–155].

One example of an interesting noise model is to assume so called depolarizing noise on each qubit

$$\rho_j \rightarrow \frac{p}{3}X_j\rho_jX_j + \frac{p}{3}Y_j\rho_jY_j + \frac{p}{3}Z_j\rho_jZ_j + (1-p)\rho_j, \quad (5.25)$$

in which case bit- and phase-flips locally correlated. In this case, the decoding of the toric code maps to computing the free energy of an eight-vertex model with quenched disorder [153].

Another interesting class of errors is so called coherent errors,

$$|\psi\rangle \rightarrow \prod_j e^{-i\phi X_j} |\psi\rangle \quad (5.26)$$

where $|\psi\rangle$ is the full wave function of the code and X_j is the Pauli- X operator acting on qubit j . These kind of errors are used to model global rotations as induced for example by stray magnetic fields or overrotations, and are relevant in platforms like trapped ions and Rydberg atoms. The mapping in this case has been derived only very recently [155] and it relates decoding to an Ising model with complex couplings. This can be reinterpreted as a 2D-Majorana scattering network and the threshold then correspond to a topological insulator-to-metal transition. The maximum-likelihood threshold in this case is found to be $p_c \approx 0.18$, quite above the threshold for incoherent errors.

A class of codes beyond surface codes that received much attention early on is color codes [156] because they allow for simpler implementations of some logical gates

[156–158]. In this case, decoding under phenomenological noise maps to computing the free-energy of a three-body Ising model for the two-dimensional code [150], and to a six-body Ising model in three dimensions [151]

Another recent result computes thresholds for fracton codes [152] and finds a threshold of $p_c \approx 0.075$ for the X-Cube code (see Sec. 2.3), which is more than twice as large as that of the three-dimensional toric code, for which $p_c \approx 0.033$.

Common to all results summarized above is that the number of logical qubits k of the codes under consideration is subextensive, that is $k/n \rightarrow 0$ as $n \rightarrow \infty$. Codes with this property are said to have “zero rate” and it turns out to be crucial for the argument presented in Sec. 5.3.2. In particular, the domain wall free energy density being nonzero only guarantees that the probability of every single logical error vanishes exponentially. This is equivalent to being below threshold only as long as the number of possible logical errors, which is 2^k , is subexponential.

In the next chapter, we present a generalization of the statistical mechanics mapping to a family of codes with finite rate, that is where $k/n \rightarrow \text{const}$ as $n \rightarrow \infty$.

6. The Statistical Mechanics of Decoding Hyperbolic Surface Codes

Parts of the work presented in this chapter have appeared in the following articles:

[159] Nikolas P. Breuckmann, Benedikt Placke, and Ananda Roy. “Critical properties of the Ising model in hyperbolic space”. In: *Phys. Rev. E* 101 (2 Feb. 2020), p. 022124. DOI: 10.1103/PhysRevE.101.022124

[160] Benedikt Placke and Nikolas P. Breuckmann. “Random-bond Ising model and its dual in hyperbolic spaces”. In: *Phys. Rev. E* 107 (2 Feb. 2023), p. 024125. DOI: 10.1103/PhysRevE.107.024125

As discussed in chapter 5, it is well known that the decoding of zero-rate stabilizer codes can be mapped onto statistical mechanics models with quenched disorder. Statistical mechanics mappings were first developed for estimating the maximum-likelihood decoding threshold of the toric code under phenomenological noise models [136]. They have since been developed for other stabilizer codes [150–152] and more general noise models [153–155]. Recently, quantum LDPC codes (that is codes with finite-weight stabilizers) with finite rate have generated much interest because they could significantly reduce the overhead required for quantum error correction. This makes a generalization of the statistical mechanics mappings to this class of codes desirable. Such a generalization is non-trivial because the mapping has to account for the fact that there is an exponential (in the number of physical qubits) number of possible logical errors. In the language of statistical mechanics this implies that they have a non-vanishing contribution to the entropy of the system. This is in contrast to the case of decoding finite-rate codes, where the number of possible logical errors does not scale with the number of physical qubits. We solve this problem for the case of hyperbolic surface codes, which are close cousins of the (euclidean) surface code but have finite rate. In particular, we show that decoding them under independent bit- and phase-flip noise maps onto what we call the “dual random-bond Ising model” in hyperbolic space. This model is the Kramers-Wannier dual of the Ising model in the absence of disorder but not identical to the Ising model even on self-dual lattices. The lack of self-duality for the Ising model in hyperbolic space is also interesting on its own, and resolves a long standing open question regarding the position of the critical point of the Ising model on regular tessellations of closed, hyperbolic manifolds [161].

The remainder of this chapter is structured as follows: after a brief introduction of hyperbolic surface codes in Sec. 6.1, we discuss the mapping and its relation to the Kramers-Wannier duality in Sec. 6.2. As an application of the mapping, in Sec. 6.2.4

we compute the maximum likelihood decoding threshold of a range of hyperbolic surface codes under independent bit- and phase-flip noise.

6.1. Hyperbolic Surface Codes

6.1.1. Hyperbolic surfaces

The hyperbolic plane is a two-dimensional manifold of constant negative curvature. It can be realized in terms of several models. Here, we will employ the *Poincaré disk model*, which is defined as follows. Consider a disk in \mathbb{R}^2 with unit radius and centered at the origin. Let x and y denote the standard coordinates of \mathbb{R}^2 . Then, the hyperbolic plane is given by the set of points

$$\mathbb{H}^2 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}, \quad (6.1)$$

with metric given by

$$ds^2 = \frac{dx^2 + dy^2}{(1 - x^2 - y^2)^2}. \quad (6.2)$$

It is immediate from Eq. (6.2) that length scales are highly distorted towards the boundary of the disk compared to the euclidean metric, see Fig. 6.1 (a,b).

Just as regular euclidean space can be tessellated by squares, triangles or hexagons, hyperbolic space can be tessellated by regular polygons as well. In fact, it turns out that hyperbolic space supports an infinite number of such regular tessellations. We label regular tessellations by the *Schläfli symbol* $\{r, s\}$, where r is the number of sides of the polygonal plaquettes and s is the number of plaquettes meeting at each vertex. For example, the hexagonal lattice has Schläfli symbol $\{6, 3\}$. Its dual lattice can be obtained by reversing the Schläfli symbol, i.e. the triangular lattice $\{3, 6\}$. These two examples, together with the self-dual square tessellation $\{4, 4\}$ are all the possible regular tessellations of the euclidean plane. The hyperbolic plane supports any regular tessellation $\{r, s\}$ as long as $1/r + 1/s < 1/2$. The $\{5, 5\}$ and $\{7, 3\}$ tessellation of the hyperbolic plane in the Poincaré disk model are shown in Fig. 6.1 (a) and (b) respectively.

In order to approximate the infinite hyperbolic plane for numerical analysis, we can consider sequences of finite neighborhoods B_R (discs) of increasing radii R . This is commonly done in the context of statistical mechanics models in euclidean space for performing finite size analysis. The models differ at the boundaries of the finite regions from the infinite euclidean plane. In the euclidean case, the effects of this deviation vanish in the thermodynamic limit as $\text{vol}(\partial B_R) / \text{vol}(B_R) \rightarrow 0$ for $R \rightarrow \infty$. This is not the case in hyperbolic space where $\text{vol}(\partial B_R)$ and $\text{vol}(B_R)$ have the same asymptotic scaling. This means that taking finite neighborhoods with boundaries can not be used to analyze the behaviour of the infinite model. We solve this problem by considering families of boundaryless, finite surfaces (supporting the same tessellation) which are indistinguishable from the infinite hyperbolic plane in local regions of increasing size at any point.

Introducing periodic boundary conditions is a much more subtle process in hyperbolic spaces compared to euclidean spaces. In particular, closed, orientable hyperbolic manifolds have a genus that is proportional to their area. This is seen




Figure 6.1.: (a, b) Poincaré disk model of the infinite hyperbolic plane \mathbb{H}^2 with the $\{5, 5\}$ and $\{7, 3\}$ lattice respectively. All edges have the same length with respect to the hyperbolic metric, see Eq. (6.2). In (b), also the surface code stabilizers are illustrated. Panel (c) shows the compactification of a finite patch of the $\{7, 3\}$ tessellation into the Klein quartic, a closed hyperbolic surface of genus $g = 3$. Faces are colored to guide they eye.

most easily by considering a theorem due to Gauß–Bonnet, which states that the geometry (curvature) of a 2D surface is connected to its topology. More concretely, it states that for any orientable surface S of genus g it holds that

$$2 - 2g = \frac{1}{2\pi} \int_S \kappa dA \quad (6.3)$$

where on the right hand side we integrate the curvature κ at every point in S over the area of S . If S is euclidean, then the curvature κ is equal to 0 at every point. From Eq. (6.3) it then immediately follows that all orientable euclidean surfaces are tori ($g = 1$). On the other hand, if S is hyperbolic then $\kappa < 0$ and constant everywhere. Orientable hyperbolic surfaces hence have

$$(-\kappa) \frac{\text{area}(S)}{2\pi} = 2g - 2 \quad (6.4)$$

so that larger surfaces necessarily have a higher genus. In Fig. 6.1 (c) we show an example of a closed $g = 3$ hyperbolic surface, called *Klein quartic*, which supports a $\{7, 3\}$ tessellation.

The curvature of the manifold supporting a given tessellation can be related to its the Schläfli symbol by considering the hyperbolic law of cosines for a triangle with internal angles (α, β, γ) and a the side-length opposing α

$$\cos(\alpha) = -\cos(\beta) \cos(\gamma) + \sin(\beta) \sin(\gamma) \cosh(\kappa a). \quad (6.5)$$

Triangulating the face of the tessellation by drawing lines from the center of the face to a vertex and the mid-point of an edge gives a triangles with angles $(\alpha, \beta, \gamma) = (\pi/r, \pi/2, \pi/s)$ and $a = 1/2$. This implies

$$\kappa = -4 \cosh^{-1} \left[\frac{\cos(\pi/r)}{\sin(\pi/s)} \right]. \quad (6.6)$$

As it turns out, the subtlety that hyperbolic surfaces are topologically complex becomes important in the Kramers–Wannier duality and also in the statistical mechanics mapping for the decoding of hyperbolic surface codes. This is because both constructions are sensitive to the number of closed loops (cycles) in the lattice and the higher genus of hyperbolic surfaces introduces more such loops.

Type	n	k	d	Type	n	k	d
$\{5, 4\}$	30	5	3	$\{5, 5\}$	15	7	3
$k/n \rightarrow \frac{1}{10}$	40	6	4	$k/n \rightarrow \frac{1}{5}$	40	10	4
	160	18	6		80	18	5
	360	38	8		150	32	6
	1710	173	9		900	182	8

 Table 6.1.: Values $[n, k, d]$ of some finite instances of hyperbolic surface codes [162].

6.1.2. Hyperbolic surface codes

Hyperbolic surface codes are close cousins of the toric code. Their microscopic construction is indeed completely analogous, and their distinct properties are fully owed to the different underlying geometry of the lattices. Consider any regular $\{r, s\}$ tessellation of the hyperbolic plane and define, just as in Eq. (5.11), one stabilizer for each vertex and face

$$A_v = \prod_{e \in \delta v} Z_e, \quad B_p = \prod_{e \in \partial p} X_e. \quad (6.7)$$

This is also illustrated in Fig. 6.1 (b), for the $\{7, 3\}$ tessellation. The stabilizers commute by the same arguments as used for the construction of the toric code and taking the code space to be the common +1 Eigenspace of all stabilizers then yields a valid code.

Introducing periodic boundary conditions amounts to putting the above construction on a closed hyperbolic manifold. As a minimal example, we show in Fig. 6.1 (c) the Klein quartic with genus $g = 3$ which supports a finite patch of the $\{7, 3\}$ tessellation. A whole family of such finite patches with periodic boundaries can be constructed [163], however the number of sites N_v for which a compactification is possible does not follow a simple pattern, in contrast to tessellations of flat space.

Asymptotically, the properties of the code family for a given tessellation is given by geometry. First, as the number of qubits n is increased, the area of the manifold tiled by the finite patch increases (for large n) proportionally to n and because of Eq. (6.4), so does its genus g . Recalling the fact that the number of qubits of a surface code is twice its genus [Eq. (5.12)], this establishes the finite rate of surface code $n \propto k$. The proportionality factor according to Eq. (6.4) is the curvature of the underlying manifold, given in terms of r and s in Eq. (6.6). The distance of hyperbolic codes is given by the size of the smallest handle of a given manifold. While this grows as the square root of n for tori, it only grows logarithmically for compact hyperbolic manifolds [164, 165]. This is the major downside of hyperbolic codes, as it implies by the arguments in Sec. 6.2.4 that errors are suppressed only with a power law in n .

To summarize, hyperbolic codes are $[n, k, d] = [n, \alpha n, \beta \log(n)]$ codes, for α, β constants depending on the specific tessellation $\{r, s\}$. Finally, note that while asymptotic properties for large n are interesting and useful, for the use in error correction in practice we are mostly interested in the properties of specific finite-size instances, the necessary size of which depends on the physical error rate. Because of this, in Table 6.1, we list the triple $[n, k, d]$ for some specific finite-size instances. Especially the $\{5, 4\}$ code for small system sizes outperforms equivalent toric codes

quite significantly. For example, four copies of the toric code with 72 qubits each (that is 288 qubits in total) are needed to encode $k = 16$ qubits at distance $d = 6$, while the $\{5, 4\}$ code provides $k = 18$ logical qubits at the same distance, using only $n = 160$ physical qubits.

6.1.3. The statistical mechanics of decoding hyperbolic surface codes

Decoding for hyperbolic code works exactly the same way as it works for the toric code. Assuming independent bit- and phase-flip noise, the task of a decoder is to match syndromes, and again minimum-weight perfect matching can be used in principle [162, 163, 166]. An interesting question is then whether MWPM is similarly close to the optimal decoder as it is for the toric code.

To answer this question, we can again map decoding to a statistical mechanics model. As already mentioned at the end of Sec. 5.3.3, an essential ingredient to the correspondence of the maximum-likelihood (ML) decoding threshold of the toric code and the critical point in the random bond Ising model is that the toric code has zero rate. For this reason we have to adjust the mapping to compute the decoding threshold of hyperbolic codes, which have finite rate.

To this end, note first that the probability of a certain error class given a syndrome [Eq. (5.14)] is the same in the hyperbolic case, and it also maps naturally to the random bond Ising model in hyperbolic space, as outlined in Eq. (5.16). In particular, recall that we can write the recovery operation proposed by the maximum-likelihood decoder as

$$R_{\text{ML}}(S) = E \circ \bar{L}_{\text{ML}}, \quad (6.8a)$$

$$\bar{L}_{\text{ML}} := \arg \max_{\bar{L}} [\bar{p}(E \circ \bar{L})], \quad (6.8b)$$

$$\bar{p}(E \circ \bar{L}) := Z_{\bar{L}}[E \hat{=} \{J_{ij}\}]. \quad (6.8c)$$

Where as in the previous section, \bar{L} denotes all possible combinations of logical operators including the identity $\bar{L} = \bar{I}$, and $Z_{\bar{L}}[E \hat{=} \{J_{ij}\}]$ is the partition function of the random-bond Ising model [Eq. (5.17)], with the couplings $\{J_{ij}\}$ set by the error E , and boundary conditions set by the logical operator \bar{L} . More specifically, we first set all couplings to be antiferromagnetic if the corresponding edge is part of the error E and ferromagnetic else

$$J_{ij} = \begin{cases} +1 & \langle ij \rangle \in E \\ -1 & \text{else.} \end{cases} \quad (6.9)$$

The boundary conditions are then set as follows: if the logical \bar{L} acts nontrivially on the qubit on edge e , we flip the sign of the coupling on the corresponding edge $J_{ij} \rightarrow -J_{ij}$, $(i, j) \hat{=} e$.

Since we have taken here without loss of generality the trial error to be the actual error ($E_0 = E$), the recovery succeeds iff the most likely logical error class is the trivial one, $\bar{L}_{\text{ML}} = \bar{I}$. This is equivalent to the free energy of the RBIM to be minimized by periodic boundary conditions. This has to happen when averaged

over all errors E , yielding the condition

$$\sum_{\bar{L} \neq \bar{I}} \sum_E \underbrace{\left(\prod_{\langle ij \rangle} p(J_{ij}) \right)}_{=: [Z_{\bar{L}}]} Z_{\bar{L}} [E \cong \{J_{ij}\}] \rightarrow 0 \quad \text{for } n \rightarrow \infty \quad (6.10)$$

where $[\bullet]$ denotes the disorder-average and

$$p(J_{ij}) = p\delta(J_{ij} - 1) + (1 - p)\delta(J_{ij} + 1). \quad (6.11)$$

Such an average over disorder is also called quenched disorder.

In flat space, the above condition is exactly equivalent to the RBIM being in its ferromagnetic phase, since a nonzero average domain-wall free energy $\langle f_{DW} \rangle > 0$ implies that $[Z_{\bar{L}}] \rightarrow 0$ for all $\bar{L} \neq \bar{I}$ and the sum is over a finite number of logical operators \bar{L} . In hyperbolic space, it is not enough for each term in the sum over logicals \bar{L} to vanish individually, since the number of logical operators is itself exponentially large in the system size. Hence, the RBIM being in its ferromagnetic phase is not sufficient for the code being below threshold.

To still obtain a mapping to a phase transition we consider a new model which we call the *dual* random-bond Ising model. The reason for this name is that, as we will see later, it is indeed the Kramers-Wannier dual of the Ising model in the absence of disorder. It is defined by augmenting the RBIM with an additional set of degrees of freedom $\{\eta_\ell\}$, encoding the boundary conditions of the model around all nontrivial loops ℓ of the *dual* lattice [see Fig. 6.2 (a)]

$$H = \sum_{\langle ij \rangle} J_{ij} \left(\prod_{\ell | \langle ij \rangle \in \ell} \eta_\ell \right) \sigma_i \sigma_j. \quad (6.12)$$

Above, the couplings J_{ij} are set by quenched disorder but the coupling on each edge $\langle ij \rangle$ also gets an additional factor, depending on the variables η_ℓ . Choosing one nontrivial loop ℓ on the dual lattice for each of the k handles of the manifold, the product in Eq. (6.12) for each edge $\langle ij \rangle$ is over all loops ℓ that include the edge (technically, its dual), which we denote by $\ell | \langle ij \rangle \in \ell$ (read “ ℓ such that $\langle ij \rangle$ in ℓ ”).

Summing in the partition function Z not only over all configurations of Ising variables $\{\sigma_j\}$ but also over all configurations $\{\eta_\ell\}$ makes the sum over boundary conditions in Eq. (6.10) explicit. Since each boundary condition corresponds to a possible logical error, this internalizes the exponential number of possible logical errors \bar{L} as an additional contribution to the entropy. Taking a loop variable to be antiferromagnetic $\eta_\ell = -1$ forces a domain wall into the system and hence has the same effects as choosing antiperiodic boundary conditions. This is costly in the ferromagnetic phase, where we expect that $\langle \eta_\ell \rangle \rightarrow 1$. Since the model as defined in Eq. (6.12) includes the finite entropic contribution of the logical errors, it being in the ferromagnetic phase thus indeed implies Eq. (6.10). The maximum-likelihood decoding threshold of hyperbolic surface codes hence corresponds not to the phase transition of the random-bond Ising model, but its *dual* in hyperbolic space.

The relation this dual model to the RBIM in terms of the Kramers-Wannier duality, as well as its phase diagram will be discussed in the following.

6.2. The Ising Model and its Dual in Hyperbolic Spaces

While the discussion above constitutes our main motivation to study the Ising model in hyperbolic space, condensed matter physics and statistical mechanics in curved spaces is interesting on its own account, and has been a subject of intense study. Curvature is known, for example, to alter the critical properties of statistical mechanics models [167], circuit quantum electrodynamics [168, 169] and band theory [170–174]. The Ising model in curved space has, to the best of our knowledge, so far only been studied in the absence of disorder [159, 161, 175–177]. In this limit, the model undergoes a phase transition from a paramagnetic high-temperature to a low-temperature ferromagnetic phase, just as its flat-space counterpart. The transition is mean-field in nature, but surprisingly it is not located at the fixed-point of the Kramers–Wannier duality, even on self-dual tessellations of the hyperbolic plane. This observation implies either the existence of a second phase transition, for which no evidence was found numerically, or a violation of self-duality of the Ising model on self-dual hyperbolic lattices. We note that the existence of a second phase transition for the disorder-free Ising model on the hyperbolic plane with free boundary condition has been proven [177–179].

Studying the Kramers–Wannier duality in the presence of curvature is interesting on its own right [180]. However, as Polyakov pointed out already in 1987 [181, Chapter 9], its understanding will have consequences also for related constructions. This includes Polyakov’s original example, the Fermionization of Ising spins [182] but as we will see in the following also the statistical mechanics mapping for decoding hyperbolic codes. Specifically, we show that the hyperbolic Ising model is not self-dual even on self-dual lattices, but instead is related by the Kramers–Wannier duality to the *dual-RBIM*, introduced in Eq. (6.12). Note that what we call the dual-RBIM it is not related to the RBIM by an exact duality in the presence of disorder.

6.2.1. Duality in the hyperbolic Ising model

We consider the Ising model (for the time being *without* quenched disorder) on a lattice $\mathcal{L} = (V, E, F)$. We denote by V the set of vertices, by E the set of edges and by F the set of faces of the lattice. Denoting nearest neighbor bonds between two vertices i and j of the lattice by $\langle ij \rangle$, the Hamiltonian of the Ising model is then given by

$$H = J \sum_{\langle ij \rangle} \sigma_i \sigma_j, \quad (6.13)$$

where $\sigma \in \{\pm 1\}$ are Ising spin variables and we assume $J < 0$ for ferromagnetic coupling.

In euclidean space, the Kramers–Wannier duality [183] relates the high-temperature expansion of the Ising model [Eq. (6.13)] to its low-temperature expansion of the same model on the dual lattice.

In particular, Kramers and Wannier showed a exact relation the two partition functions

$$Z(T) = \tilde{Z}(T^*) \quad (6.14a)$$




Figure 6.2.: Duality on the Klein quartic, with the $\{7, 3\}$ tessellation shown in solid lines, and its dual $\{3, 7\}$ tessellation indicated by dashed lines. In (a) we show a nontrivial loop (that is a cycle that is not a boundary) on the dual lattice in blue, and its corresponding cocycle on the primal lattice in red. In (b), we show trivial loop (that is a boundary) on the primal lattice in blue and its corresponding coboundary on the dual lattice in red.

where Z and \tilde{Z} are the partition functions of the Ising model on the lattice and its dual respectively, and T and T^* satisfy

$$\sinh(2J/T) \sinh(2J/T^*) = 1. \quad (6.14b)$$

On a self-dual lattice, $Z = \tilde{Z}$ and thus the duality [Eq. (6.14)] constitutes an exact mapping between the behavior of the system at high and low temperature. In particular, assuming that a single phase transition occurs, this fixes the critical temperature to the fixed-point of Eq. (6.14b)

$$\sinh(2J/T_c) \sinh(2J/T_c) = 1 \Rightarrow T_c \approx 2.2692J. \quad (6.15)$$

An open question posed by earlier studies [159, 161] was how Eq. (6.15) is violated in hyperbolic space. That is, if the Kramers–Wannier duality [Eq. (6.14)] applies also to hyperbolic lattices, one of the following must hold: either all self-dual hyperbolic lattices (that is tessellations of compact hyperbolic manifolds with Schläfli symbol $\{r, s\}$ with $r = s$) have the *same* critical temperature, given by Eq. (6.15), or there exist *two* phase transitions, related by Eq. (6.14b). In fact, as we will show below, the Ising model on tessellations of compact hyperbolic manifolds is not related by the Kramers–Wannier duality to the same Ising model on the dual lattice. In particular, it is not self-dual, even on self-dual tessellations.

Re-derivation of the Kramers–Wannier duality

To understand this, let us perform a careful re-derivation of the Kramers–Wannier duality. To this end, we first consider the high-temperature expansion of the Ising model on a lattice $\mathcal{L} = (V, E, F)$. Let Z_1 be the set of subsets $\gamma \subset E$ such that in the subgraph induced by any such γ , every vertex has even degree. The subsets $\gamma \in Z_1$ are called *cycles*. It is well-known that the partition function can be written

as a sum over the set of all cycles of the graph (see e.g. [184, Chapter 2]):

$$Z(K) = \sum_{\sigma \in \{\pm 1\}^N} \prod_{(i,j) \in E} \exp(K\sigma_i\sigma_j) \quad (6.16a)$$

$$= (\cosh K)^{|E|} \sum_{\sigma} \prod_{(i,j)} (1 + \sigma_i\sigma_j \tanh K) \quad (6.16b)$$

$$= 2^N (\cosh K)^{|E|} \sum_{\gamma \in Z_1} (\tanh K)^{|\gamma|} \quad (6.16c)$$

where we have defined $K = -J/T$ and $|S|$ denotes the size of the set S .

Note that the set Z_1 of cycles γ in Eq. (6.16) includes ones that are contractible as well as ones that are non-contractible. Two examples for such cycles, on a surface with genus 3, tessellated by the $\{7, 3\}$ tessellation (cf. also Fig. 6.1), are given in Fig. 6.2. On the right we show a contractible cycle on the primal lattice (solid lines) in blue. On the left we show, also in blue, a non-contractible cycle on the dual lattice (dashed lines).

To establish the duality, we also consider the low-temperature expansion of the Ising model, but on the dual lattice $\mathcal{L}^* = (V^*, E^*, F^*) = (F, E, V)$. For regular tessellations of hyperbolic surfaces, the dual lattice is just obtained by swapping the first and second entry of its Schläfli symbol $\{r, s\}$. This is also indicated in Fig. 6.2. The primal lattice (solid lines) is the $\{7, 3\}$ tessellation and its dual (dashed lines) is the $\{3, 7\}$ tessellation of the same surface.

The low temperature expansion follows from expressing the partition function in terms of excitations on top of the (ferromagnetic) ground state. These are given by domain walls. For example, consider starting from a all-ferromagnetic state of the Ising model [Eq. (6.13)] on the dual lattice indicated by dashed lines in Fig. 6.2. The cost of flipping the spin on the central cite is given by the size of the domain wall indicated in red on the right of Fig. 6.2. Generally, let B^{1*} be the set of all possible domain walls on the dual lattice. We can write

$$\tilde{Z}(K) = \sum_{\sigma \in \{\pm 1\}^N} \prod_{(i,j) \in E^*} \exp(K\sigma_i\sigma_j) \quad (6.17a)$$

$$= 2 \sum_{\omega^* \in B^{1*}} \exp(K^*)^{|E^*| - 2|\omega^*|} \quad (6.17b)$$

$$= 2 \exp(K)^{|E^*|} \sum_{\omega^* \in B^{1*}} \exp(-2K)^{|\omega^*|} \quad (6.17c)$$

where the second equality directly follows from the definition of B^{1*} . In the language of homology, the set B^{1*} is given exactly by the set of *coboundaries* on the *dual* lattice.

The basis of the Kramers–Wannier duality, homologically speaking, is the fact that the set of cycles Z_1 is in one-to-one correspondence with the set of *cocycles* Z^{1*} on the dual lattice \mathcal{L}^* . This is also indicated in Fig. 6.2 where we show two examples of the correspondence of cocycles (red) and cycles (blue). The left side shows a non-contractible cocycle on the primal lattice (solid, red) and the corresponding cycle on the dual (blue, dashed). The right side shows a contractible cycle (a *boundary*) on the primal lattice (blue, solid) and the corresponding cocycle (a *coboundary*) on the dual lattice (red, dashed).

Using this equivalence, $Z_1 = Z^{1*}$, as well as Eq. (6.14b), and defining $K^* = J/T^*$, we can then rewrite

$$Z(K) = 2 \exp(K^*)^{|E^*|} \sum_{\gamma^* \in Z^{1*}} \exp(-2K^*)^{|\gamma^*|}. \quad (6.18)$$

Above, the right hand side is *almost* the low-temperature expansion of the Ising model on the dual lattice [Eq. (6.17)], at temperature T^* [Eq. (6.14b)]. The difference between Eq. (6.18) and Eq. (6.17) is that the sum above is over all cocycles $\gamma^* \in Z^{1*}$ whereas the low-temperature expansion is a sum over domain walls $\omega^* \in B^{1*}$, that is coboundaries or “contractible” cocycles. Physically, we can rationalize this difference by looking at the example of a non-contractible cocycle on the left of Fig. 6.2 (red, solid). The corresponding cycle (blue, dashed) appears in the high-temperature expansion of the dual lattice (every vertex in it has even degree). However, there is no set of spins on vertices of the primal lattice that we could flip to get a domain of that form.

Hence, for Ising models on regular tessellations of closed manifolds, we have established what is the *difference* between their high-temperature expansion [Eq. (6.16)] and the low-temperature expansion of the same model on the dual lattice at the dual temperature [Eq. (6.18)]. In the following, we show that (i) for tessellations of closed euclidean surfaces (tori), this difference vanishes in the thermodynamic limit, yielding the Kramers–Wannier duality [Eq. (6.14)], and (ii) the difference does *not* vanish for tessellations of closed hyperbolic surfaces, leading to a violation of Eq. (6.14).

Note that the contribution of any cocycle in Eq. (6.18) has a weight $\exp(-2K^*)^{|\gamma^*|}$. For euclidean lattices on an $L \times L$ torus, this implies that the contribution of any non-contractible cocycle is at least of order $\mathcal{O}\left(\exp(-2K^*)^L\right)$. Focussing on such minimal-size cocycles, of which there are $\sim L$, the difference between Eq. (6.18) and the low-temperature expansion of the Ising model vanishes in the thermodynamic limit

$$Z(T) - \tilde{Z}(T^*) \sim L \exp(-2K^*L) \xrightarrow{L \rightarrow \infty} 0. \quad (6.19)$$

This then yields Eq. (6.14).

In contrast, in hyperbolic space, the number of minimal, non-contractible cocycles goes as $\sim N$ [see Eq. (5.12) and discussion there] while their length grows only logarithmically [164, 165]. This means that the same difference goes as

$$Z(T) - \tilde{Z}(T^*) \sim N^{1-2K^*} \quad (6.20)$$

which does not generally vanish as $N \rightarrow \infty$.

The dual Ising model in hyperbolic space

In order to obtain a model that does fulfill the Kramers–Wannier duality, we have to define a model where possible domain walls on top of the ferromagnetic ground state include all non-contractible cocycles.

We achieve this by a rather simple trick. Given an Ising model [Eq. (6.13)] on a tessellation of a closed hyperbolic surface S with $2g$ nonequivalent, non-contractible cocycles ℓ , we introduce one additional Ising degree of freedom η_ℓ per nonequivalent,

non-contractible cocycle. This yields exactly the same model as Eq. (6.12), but without quenched disorder

$$H = J \sum_{\langle ij \rangle} \left(\prod_{\ell | \langle ij \rangle \in \ell} \eta_\ell \right) \sigma_i \sigma_j \quad (6.21)$$

where $J < 0$ as before is chosen to be ferromagnetic and we have chosen one representative per nontrivial cocycle ℓ , as illustrated in Fig. 6.2 (a). Recall that the effect of flipping this Ising degree of freedom $\eta_\ell \rightarrow -\eta_\ell$ is to force a domain wall wrapping around the nontrivial cycle ℓ . Because of this, excitations of the model defined by Eq. (6.21) include the nontrivial cocycles of the lattice and its partition is given exactly by Eq. (6.18). This establishes that in the pure limit (without disorder), what we called the dual RBIM is indeed the Kramers-Wannier dual of the Ising model.

This model also gives another rationale for the difference between duality of Ising models on tessellations of euclidian and hyperbolic manifolds. Strictly speaking, the Kramers-Wannier dual of the Ising model on *finite* tessellations of euclidean manifolds is also given by Eq. (6.21). However since all closed, orientable euclidean manifolds are tori, the dual model has only two additional degrees of freedom ℓ compared to the original Ising model [Eq. (6.13)]. Hence, they have no finite entropic contribution in the thermodynamic limit and the dual model has the same thermodynamic properties as the original model. In contrast, on tessellations of closed hyperbolic manifolds, the number of additional variables ℓ in Eq. (6.21) is extensive ($\sim N$) and hence changes the properties of the model, even in the thermodynamic limit.

6.2.2. The random-bond Ising model

The random-bond Ising model (RBIM), first introduced by Edwards and Anderson [41] to model the interaction of dilute magnetic alloys, serves as a simple model to study critical phenomena in systems with quenched disorder. The Hamiltonian for the RBIM on a lattice with nearest-neighbor bonds $\langle ij \rangle$ as already given in Eq. (5.17) is

$$H = \sum_{\langle ij \rangle} J_{ij} \sigma_i \sigma_j \quad (6.22)$$

where $\sigma_i \in \{\pm 1\}$ are Ising spin variables and J_{ij} are random couplings. Whenever we refer to the Ising model in “hyperbolic space” or on “hyperbolic lattices” throughout this work, we refer to a model where spins are located on the vertices of regular tessellations of *compact* hyperbolic manifolds, with Schläfli symbol $\{r, s\}$. This emphasis is important, since considering the same model on non-compact hyperbolic manifolds with, for example, open or closed boundary conditions will generally change its properties [178, 179]. The couplings are distributed independently and identically according to Eq. (6.11), which is called the “ $\pm J$ -distribution”

$$P(J_{ij}) = p \delta(J_{ij} - 1) + (1 - p) \delta(J_{ij} + 1). \quad (6.23)$$

The free energy of the model, when considering quenched disorder is then given

by

$$F = [\log(Z)], \quad (6.24)$$

$$Z = \sum_{\{\sigma\}} \exp \left(-\beta \sum_{\langle i,j \rangle} J_{ij} \sigma_i \sigma_j \right), \quad (6.25)$$

where brackets $[\dots]$ denote the average over disorder configurations.

For $p = 0$, the model reduces to the ferromagnetic Ising model, which we have studied in Ref. [159]. This model as a function of temperature undergoes a phase transition from a high-temperature paramagnetic into a low-temperature ferromagnetic phase. Our study revealed that this transition is mean-field in nature for all investigated tessellations. Here, we extend this previous work to the case of finite $0 \leq p \leq 1/2$.

We also study the *dual* Ising model Eq. (6.21) in the presence of quenched disorder, as already introduced in Eq. (6.12)

$$H = \sum_{\langle ij \rangle} J_{ij} \left(\prod_{\ell | \langle ij \rangle \in \ell} \eta_\ell \right) \sigma_i \sigma_j. \quad (6.26)$$

As before, the $\sigma_j \in \{\pm 1\}$ are Ising variables, as are the $\eta_\ell \in \{\pm 1\}$. While the σ_j are located on the vertices of the lattice, each η_ℓ is associated with a nontrivial cocycle ℓ (cf. Sec. 6.2.1). The J_{ij} are random couplings drawn from the $\pm J$ distribution defined in Eq. (6.11).

The Kramers–Wannier duality [Eq. (6.14)], as usual, is only valid in the disorder-free case. However there is a conjecture by Takeda and Nishimori [185] relating the location of the Nishimori point of the RBIM with the position in the dual model

$$H(p_N) + H(p_N^*) = 1 \quad (6.27)$$

where $H(p) = -p \log_2(p) - (1-p) \log_2(1-p)$ is the binary entropy. As discussed in Sec. 6.2.4, we find that the conjecture holds approximately, but not within error bars.

Possible phases and order parameters

At high temperature, both the RBIM and its dual are in the paramagnetic phase. As the temperature is lowered, at low disorder this gives way to a ferromagnetic phase which is continuously connected to that of the pure model at $p = 0$. The transition from the paramagnet to the ferromagnet corresponds to an instability of the mean of the magnetization distribution $\rho(m)$. That means while in the paramagnet we find in the thermodynamic limit

$$\rho(m) = \delta(m), \quad (6.28)$$

in the ferromagnetic phase

$$\rho(m) = \delta(|m| - M), \quad (6.29)$$

for some finite magnetization M .




Figure 6.3.: Schematic phase diagram of the random bond Ising model and its dual on the hyperbolic plane as a function of temperature T and the fraction of antiferromagnetic bonds p . The Nishimori line is indicated in dashed-gray. The high-temperature paramagnetic phase at low temperatures gives way either to a ferromagnetic phase at weak or to a spin glass phase at strong disorder respectively. For the dual model, we indicate only the schematic boundary of the FM phase by a dash-dotted line. The critical temperatures of the disorder free models, T_c and T_c^* , are related by the Kramers–Wannier relation. The phase boundary of the dual model corresponds to the decoding threshold of the hyperbolic surface code under phenomenological noise. In particular, the critical disorder at zero temperature, $p_{c,0}^*$, and along the Nishimori line, p_N^* , are equal to the threshold of the minimum-weight perfect matching and the maximum-likelihood decoder, respectively.

For large disorder, $p \approx 1/2$, random systems can also develop spin glass order at low temperature, which corresponds to an instability in the variance of the magnetization distribution, which is also called the Edwards-Anderson (EA) order parameter

$$q_{\text{EA}} = [m^2], \quad (6.30)$$

where $[\bullet]$ as before denotes disorder averaging and the magnetization vanishes ($M = [m] = 0$).

The schematic phase diagram of the RBIM and its dual on the hyperbolic plane is shown in Fig. 6.3, with the Nishimori line [Eq. (5.18)] indicated as a gray-dashed line. For the dual model, we only indicate the phase boundary of the ferromagnetic phase. There could exist a spin-glass phase also in the dual model, but the investigation of that is beyond the scope of this work.

It can be shown exactly, that the phase boundary of any magnetized phase must be reentrant or vertical, that is no magnetized phase can exist for $p_N < p$ [141]. As discussed in Sec. 6.2.4, this is equivalent to the fact that the minimum-weight perfect matching decoder cannot have a higher threshold than the optimal decoder. In all known cases, the multicritical point in the RBIM lies on the Nishimori line, and as we will show below this is also the case in hyperbolic space.

The ferromagnetic phase of the RBIM is expected to have a larger extent than that of its dual, since the additional cocycle degrees of freedom η_ℓ have a finite contribution to the entropy, which is then strictly greater than that of the RBIM.

6.2.3. Numerical phase diagram of the RBIM

We use high-temperature series expansion of the susceptibility (HTSE) as well as Monte-Carlo simulations to map out the phase diagram of the random-bond Ising model on the $\{5, 5\}$ lattice in detail. The HTSE converges rapidly in hyperbolic space, and allows for highly accurate estimates of both the position of the critical point as well as the critical exponents (see App. C for details on both derivation of the series as well as its analysis). Naturally however, it only has access to the high-temperature regime, and cannot probe the phase boundary between the spin glass and the ferromagnetic phase. Because of this, and to corroborate the HTSE results, we also use Monte-Carlo simulations.

The phase diagram of the RBIM on the $\{5, 5\}$ tessellation, as obtained from HTSE and MC simulations is shown in Fig. 6.4. Compared to the RBIM on the euclidean square ($\{4, 4\}$) lattice, we find a much larger ferromagnetic phase and a extended spin glass phase. In contrast to the Bethe lattice, here we do not find evidence for a magnetized spin glass phase [186, 187], although our low-temperature data is not conclusive enough to fully rule out its existence.

Turning to explain our results in more detail, in Fig. 6.4 we show both, the magnetization m in panel (a) as well as the Edwards-Anderson order parameter q in panel (b) as obtained from a MC simulation with system size $N_v = 1920$. While the magnetization is nonzero only in the ferromagnetic phase, the EA order parameter is nonzero in both the ferromagnet and the spin glass. We superimpose these plots with the critical points obtained using finite-size scaling of the MC data (open circles) and with the critical lines obtained from HTSE of the (EA-) susceptibility (solid lines). In both methods, we can distinguish the transition from the paramagnet to a ferromagnetic phase and that to a spin glass phase reliably. In the finite size




Figure 6.4.: Phase diagram of the RBIM on the $\{5, 5\}$ tessellation as a function of temperature T and disorder strength p . We show both the magnetization m in panel (a) as well as the Edwards-Anderson order parameter q_{EA} in panel(b), both obtained from Monte-Carlo (MC) simulations of a $N = 1920$ system. We superimpose this with the phase boundaries obtained from the high-temperature series expansion (HTSE) and MC (see main text for details).

analysis of the MC data, a transition to the ferromagnetic phase is signaled by a crossing of both the binder cumulant of the magnetization, g as well as a crossing of the binder cumulant of the Edwards-Anderson order parameter, g_{EA} . In contrast, at the transition to a spin glass phase, only g_{EA} shows a crossing while g does not, since the magnetization m vanishes in the spin glass. Finite size scaling along the Nishimori line indicates a transition at $p_N = 0.247 \pm 0.02$. Finite size analysis as a function of temperature at constant disorder shows a transition into a ferromagnet for $p \lesssim p_N$ and a transition into a spin glass for $p \gtrsim p_N$, making the Nishimori point the multicritical point.

This result is corroborated by HTSE analysis. Here, a transition to the ferromagnet (spin glass) is signaled by the divergence ferromagnetic (EA-) susceptibility $\chi_{(EA)}$. Note that since the non-divergent susceptibility at both transitions typically also has a weak singularity (a cusp), series analysis normally predicts a divergence for both susceptibilities, but at different critical temperatures. In practice, we distinguish the two transitions by the fact which susceptibility is predicted to diverge at higher temperature. Along the Nishimori line, the two susceptibilities are equal and HTSE yields a critical point $p_N = 0.246793 \pm 4.2 \times 10^{-6}$. For $p < p_c$ we find a transition to a ferromagnetic phase at low temperature, while for $p > p_c$ we find a transition into a spin glass phase, again suggesting that the Nishimori point is indeed the multicritical point of the model.

Critical behaviour

In Table 6.2, we show results for the critical exponents of the model on the $\{5, 5\}$ tessellation for different scaling axis (with the $p = 0$ results taken from Ref. [159]). The best results are typically obtained from the HTSE. To extract estimates for

	$p = 0$	Nishimori Line	$p = 1/2$
μ	2	3.0 ± 0.1	2.0 ± 0.1
γ	1.000001 ± 0.000005	1.0003 ± 0.0008	-
γ_{EA}	-	1.0003 ± 0.0008	1.0011 ± 0.0025
β	0.46 ± 0.05	1.00 ± 0.05	-

Table 6.2.: Critical exponents on the $\{5, 5\}$ lattice along different scaling axes. We estimate the correlation volume exponent, μ , from finite size analysis of the binder parameter g . For the susceptibility exponents γ and γ_{EA} the best estimates are obtained via HTSE analysis.

the exponent also from the Monte-Carlo simulations, we have to perform finite size scaling. This is slightly nontrivial because the closed hyperbolic surfaces considered do not have a unique linear dimension.

Instead of defining a correlation length, we perform finite size scaling as a function of the number of sites N . This was initially proposed for a fully connected model [188] and has been used for hyperbolic lattices with open boundary conditions [189] as well as in our earlier study of the pure Ising model on regular tessellation of closed hyperbolic surfaces [159]. The main idea is that a quantity A , close to criticality, follows a scaling form

$$A \sim |T - T_c|^a F(N/N_c) \quad (6.31)$$

with a correlation number N_c . Assuming that a corresponding system of finite dimension $d = d_c$, where d_c is the upper critical dimension, has the same scaling behavior as its hyperbolic sibling, it follows that

$$N_c \sim |T - T_c|^{-\mu}, \quad (6.32)$$

with the critical exponent

$$\mu = \nu_{\text{MF}} d_c, \quad (6.33)$$

where ν_{MF} is the mean-field value of the critical exponent of the correlation length ξ . We use this mainly to extract the exponent of the correlation volume μ . The best finite-size scaling collapse of the Monte Carlo data along the Nishimori line is shown in Fig. 6.5. The best collapse is obtained for slightly different values of p_c for the susceptibility and the binder cumulant, which we attribute to finite size effects.

A reasonable estimate of the critical disorder along the Nishimori line is $p_N = 0.247 \pm 0.002$. Thus, the ferromagnetic order in hyperbolic space is quite a bit more robust than in flat space, where $p_N \approx 0.11$. Unfortunately, the converse behavior is expected for the threshold of the hyperbolic $\{5, 5\}$ code, which is given by the critical disorder of the *dual* model, p_N^* (cf. Fig. 6.4).

The results in Table 6.2 are all compatible with the mean-field expectation, except for the exponents $\mu = 3$ and $\beta = 1$, observed along the Nishimori line. This is because as established above, the Nishimori line passes through the multicritical point which generally shows distinct critical behavior even in (effectively) infinite dimensions. Note that still, the exponents are consistent with the hyperscaling relation

$$\mu = 2\beta + \gamma \quad (6.34)$$

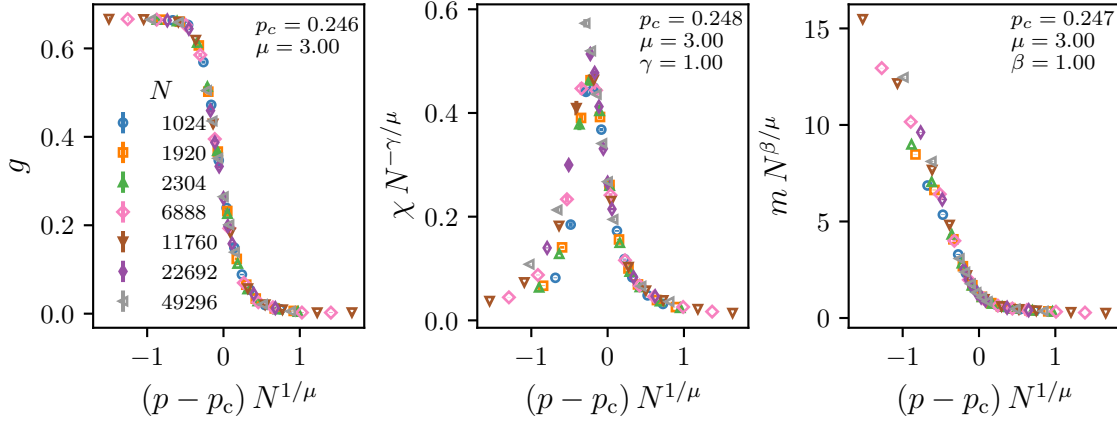


Figure 6.5.: Finite size scaling collapse of the Binder cumulant g , the susceptibility χ , and the magnetization m of the random bond Ising model on the $\{5, 5\}$ tessellation along the Nishimori line.

The specific heat does not develop a power-law singularity for any of the transitions considered and hence we do not present a critical exponent α .

In the next section we turn to the behavior of the *dual* RBIM, the phase transition of which along the Nishimori line maps onto the maximum-likelihood decoding threshold of the code.

6.2.4. Numerical phase diagram of the dual RBIM and threshold values

Disorder-free limit: Kramers-Wannier duality

In Fig. 6.6, we show results from Monte-Carlo simulations of the dual Ising model, that is Eq. (6.26) with $p = 0$, on the $\{5, 5\}$ tessellation. We show the average vertex magnetization $m = \langle \sigma_j \rangle$ and *loop* magnetization $m_\eta = \langle \eta_\ell \rangle$, as well as the Binder cumulants g and g_η for vertex and loop magnetization, respectively. The fact that the magnetizations for different system sizes cross at a single point, together with the pronounced dip of the Binder cumulants just before the transition, provides evidence that both quantities undergo a strongly first-order transition. This result is rather surprising, given that the magnetization of the Ising model on the same lattice undergoes an ordinary second-order transition and the two models are related by the Kramers-Wannier duality. While the numerical results are quite suggestive, further analysis would be needed to make a definitive statement about the order of the transition given the very small system sizes accessible.

Since the $\{5, 5\}$ lattice is self dual and the Kramers-Wannier duality [Eq. (6.14b)] is exact at $p = 0$, we expect the transition to occur at a critical temperature dual to the the critical point of the Ising model. Substituting $T_c = 3.93$ [159] into Eq. (6.14b) yields $T_c^* \approx 1.44$, which we indicate in Fig. 6.6 by a vertical dashed line and is in good agreement with the position of the crossing of both Binder cumulants and magnetizations.

To corroborate the above findings, we also implement the Wang-Landau algorithm [190–193] and compute the free energy difference of the Ising model and its dual that




Figure 6.6.: Evidence for a phase transition at T_c^* of the pure dual Ising model (that is Eq. (6.26) with $p = 0$) on the $\{5, 5\}$ lattice. We show the vertex magnetization $m = \langle \sigma_j \rangle$, its Binder cumulant g as well as the loop magnetization $m_\eta = \langle \eta_\ell \rangle$ and its Binder cumulant g_η as a function of temperature.

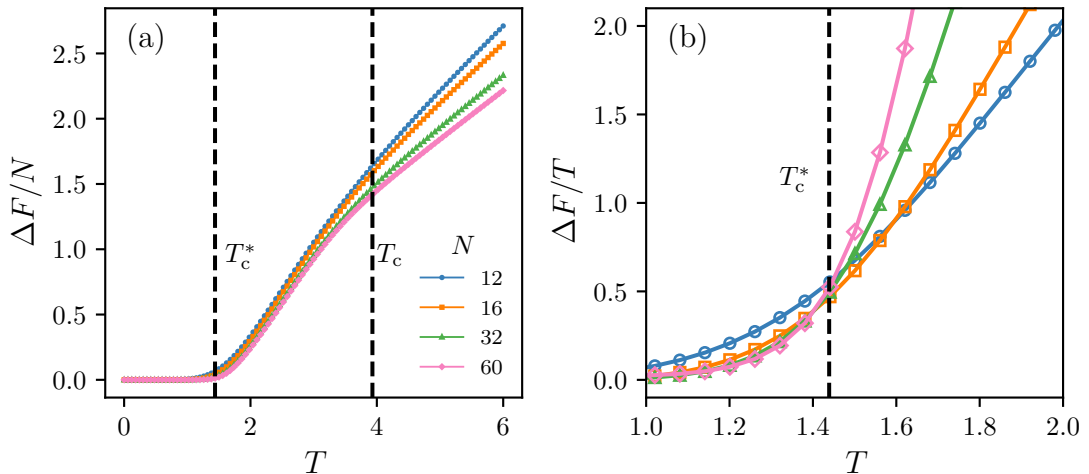


Figure 6.7.: Free energy difference [Eq. (6.35)] between the Ising model and its dual on the $\{5, 5\}$ lattice. Panel (a) shows the difference and per site. Plotting the total free energy difference per temperature in (b) yields another way to estimate the critical temperature.




Figure 6.8.: Free energy difference [Eq. (6.35)] between the random bond Ising model (RBIM) and the dual-RBIM on the $\{5, 5\}$ lattice along the Nishimori line. Solid lines are cubic fits to the data. The shaded region indicates the estimated location of the Nishimori point $p_N = 0.0222 \pm 0.0020$. The inset shows the best data collapse, assuming the same correlation exponent μ as in the RBIM.

is

$$\Delta F(T) = T (\log[Z_{\text{tot}}(T)] - T \log[Z_0(T)]) . \quad (6.35)$$

Here, Z_{tot} is the partition function of the dual Ising model, that is it includes a sum over all cocycle variables η_ℓ (therefore the subscript ‘tot’). Z_0 is the partition function of the Ising model on the same lattice, that is we fix $\eta_\ell = 1$ for all ℓ . Because of the latter relation between Z_{tot} and Z_0 , we have $\Delta F > 0$ for all T . In the ordered phase of the dual Ising model, the difference vanishes since the sum over cocycle variables does not contribute. This is shown in Fig. 6.7. The quantity ΔF also has the advantage of indicating both phase transitions in one observable, since the free energy of the Ising model shows a visible kink at T_c . Both critical temperatures are again indicated in the figure by vertical dashed lines.

Dual random-bond Ising model

In the case of the random model, the dual-RBIM is not exactly dual to the RBIM and hence we have no a-priori guess for the location of the critical point. Additionally, as is already the case in the pure model, the strongly first order nature of the transition complicates its numerical investigation. We find that single-spin flip Monte-Carlo is unreliable even for small system sizes. Since the Wang-Landau algorithm performs a random-walk in phase space, it is less sensitive to large energy barriers (instead, it is sensible to entropic bottlenecks). Most importantly, the flatness of the entropy histogram $S(m)$, together with the known extremal values of the magnetization $m \in [-1, 1]$ can be used to reliable test the convergence of the WL algorithm. We find that it is still converging up to system sizes $N_v = 60$ and we can hence get a (admittedly rather rough) estimate of the location of the critical point from the free

energy difference [Eq. (6.35)]. The difference for $T < T_c^*$ vanishes as a function of system size and diverges as a function of system size for $T_c^* < T$ respectively. In Fig. 6.8, we show ΔF as a function of disorder strength p along the Nishimori line [Eq. (5.18)]. The data is consistent with a transition at $p_N^* = 0.0222 \pm 0.0020$, which is indicated in the figure by a shaded area. The inset shows the best data collapse assuming the same correlation exponent $\mu = 3$ as in the RBIM.

Substituting the value of $p_N = 0.246793 \pm 4.2 \times 10^{-6}$ obtained from the high-temperature expansion of the RBIM into the duality relation conjectured by Nishimori (Eq. (6.27)) and solving for p_N^* yields a value of $p_N^* = 0.029891 \pm 2 \times 10^{-6}$. As observed for the RBIM on a range of euclidean lattice geometries [185] this is somewhat close to our numerical result but not compatible within error bars.

The above results yields, as discussed before, the maximum-likelihood decoding threshold of the hyperbolic $\{5, 5\}$ code

$$p_{\text{th,ML}} = p_N^* = 0.0222 \pm 0.0020. \quad (6.36)$$

This can be compared to the known threshold when using a minimum-weight perfect-matching decoder $p_{\text{th,MWPM}} \approx 0.0175$ [162]. Using an optimal decoder rather than MWPM hence increases the threshold by about 27%.

6.3. Conclusion and Outlook

To summarize, the results presented in this chapter answer two open questions from quite distinct areas of physics. First, we generalized for the first time the well-known statistical mechanics mapping of decoding zero rate codes to a family of codes with finite rate. In particular, we have established that the maximum-likelihood decoding threshold of hyperbolic surface codes corresponds to a paramagnet to ferromagnet transition along the Nishimori line of a novel model, which we call the dual random-bond Ising model. We also showed that this model is in the disorder-free limit the Kramers-Wannier dual of the Ising model in hyperbolic space. This settles a second open question, regarding the position of the critical point of the hyperbolic Ising model away from the fixed point of the Kramers-Wannier duality [159, 161]. The solution to this conundrum is simply that the Ising model on tessellations of closed hyperbolic manifolds is not self dual, even on self-dual tessellations.

Motivated by these results, we have then conducted an in-depth numerical study of the random bond Ising model (RBIM) as well as its dual in hyperbolic space. Combining high-temperature expansion techniques and Monte-Carlo simulations, we mapped out the phase diagrams of both models. Studying the critical properties of the high-temperature transitions, we showed that with the exception of the multicritical Nishimori point, all transitions are mean-field in nature. We verified the duality of both models explicitly in the disorder-free case and showed that the extended duality as conjectured by Takeda, Sesamoto and Nishimori [185] is fulfilled only approximately. Finally, we computed the critical value of disorder of the dual RBIM along the Nishimori line, which yields the first estimate of the maximum-likelihood decoding threshold of the $\{5, 5\}$ hyperbolic surface code under independent bit- and phase-flip noise, $p_{\text{th}} = 0.0222 \pm 0.0020$.

This work opens up multiple interesting areas for future work. For example, beyond the scope of the current study was a detailed investigation of the nature of

the spin-glass phase in hyperbolic space and its fate in the dual-RBIM. Moreover, a detailed investigation of the phase space structure of the dual model could yield valuable insights into the decoding of finite-rate quantum codes.

Part III.
Materials

7. The Hunt for Quantum Spin Ice

Frustrated magnets consisting of corner sharing fully connected clusters are known to host a wealth of unusual phenomena. Frustration describes the inability to simultaneously minimize the energy of all magnetic interactions. It suppresses the tendency to order, which is the usual fate of interacting magnets at low temperature. One fascinating possibility — theoretically proposed by Philip Anderson 50 years before the writing of this thesis [12] — is that quantum fluctuations enable a system to avoid this fate even at zero temperature. The resulting state is called a quantum spin liquid (QSL). While QSLs are well established theoretically to exist in many model Hamiltonians [17, 25, 194–196], they still await conclusive observation in experiment, although several candidate materials showing some evidence of QSL behavior exist [197].

We begin this chapter by introducing a model called spin ice, which in its classical regime is known to be realized in a range of materials. This classical regime already hosts fractionalized excitations and exotic dynamics. Moreover, spin ice also has a quantum counterpart — quantum spin ice — which is one of the most promising candidates for a QSL that could also be realized in a material. This is because it is one of the very few models for which direct material realization exists and whose quantum disordered ground state is also well established in theory.

One realization of a particularly simple Hamiltonian known to host a quantum spin ice ground state is provided by so called dipolar-octupolar pyrochlores. This family of magnetic materials will be the subject of the second half of this chapter, where we discuss their general pseudospin Hamiltonian as well as the slightly peculiar relation of these pseudospins to the local magnetic moment.

7.1. Classical Spin Ice

The simplest model describing a classical spin ice phase (again, first proposed by Anderson [198]) is given by the antiferromagnetic Ising model on the pyrochlore lattice [Fig. 7.1 (a)]. In this case, the Hamiltonian can be rewritten as a sum over tetrahedra, in the same way as discussed in Sec. 2.1.2 for the Heisenberg model

$$H = J \sum_{\langle ij \rangle} \sigma_i \sigma_j \quad (7.1a)$$

$$= \frac{J}{2} \sum_t M_t^2 + \text{const}, \quad (7.1b)$$

$$M_t = \sum_{j \in t} \sigma_j. \quad (7.1c)$$

Above, the sum in the first line is over all nearest-neighbor bonds $\langle ij \rangle$, and the sum in the second line is over all tetrahedra in the lattice t . Since the latter is (up to a

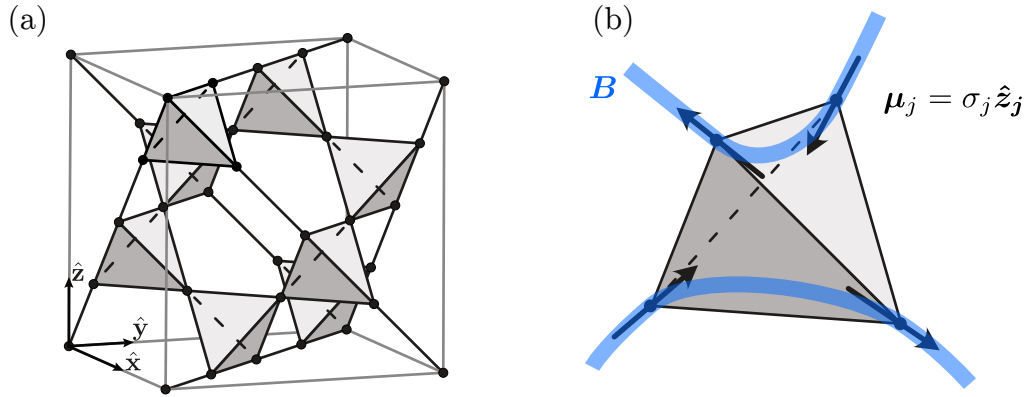


Figure 7.1.: The pyrochlore lattice is formed by a network of corner sharing tetrahedra as shown in (a). In Spin Ice, local magnetic moments (black arrows) fulfill a local two-in-two out rule on each tetrahedron, which can be interpreted as a Gauss law for an emergent magnetic field (blue).

constant) a sum over nonnegative terms, the ground state is given by the condition $M_t = 0 \forall t$.

7.1.1. The ice rule and emergent Maxwell theory

This constraint is also called the “ice rule” because an equivalent rule describes proton disorder in water ice, as first studied by Bernal and Fowler in 1933 [199]. In cubic water ice (ice I_c)¹, oxygen molecules are arranged on a diamond lattice, which is given exactly by the centerpoints of pyrochlore tetrahedra. Each oxygen is connected to its four neighbors by a hydrogen bond, which is however asymmetric, with the shared proton always being a bit closer to one of the bonded partners. The collective energy is minimized by configurations equivalent to those discussed for the magnetic model above, where each oxygen has two close and the two far protons. Pauling in 1935 [200] gave a remarkably accurate estimate of the residual entropy, that is the logarithm of the number of ground states, by treating the constraints as independent on each tetrahedra. Since 6 out of 16 states fulfill the ice rule on every tetrahedron, the total number states in the ground state is smaller than the total number of states, 2^N , by a factor of 6/16 for each of the $N_t = N/2$ tetrahedra, yielding

$$S_0 = k_B \log \left[2^N \left(\frac{6}{16} \right)^{N_t} \right] = k_B \frac{N}{2} \log \left(\frac{3}{2} \right), \quad (7.2)$$

where N is the total number of lattice sites. Including correlations between the constraints on different tetrahedra changes the above estimate by less than two percent [201, 202].

As we will see in the following, the ice rules can be interpreted as a Gauss’s law an an emergent magnetic field \mathbf{B} . In spin ice, this field is proportional to the magnetization density. This is because the local magnetic moments are given as

$$\boldsymbol{\mu}_j = \sigma_j \hat{\mathbf{z}}_j, \quad (7.3)$$

¹The more common form, that is hexagonal water ice (ice I_h), obeys the same rules but on a different lattice. It also has the same residual entropy.




Figure 7.2.: Structure factor $S(\mathbf{q})$ of spin ice [Eq. (7.1)] at low temperature $T = 10^{-3}J$. In addition to Eq. (7.5), the calculation includes the dipolar form factor (see Sec. 8.1 for details). Pinch-point singularities appear at zone centers $\mathbf{q} = [111]$ and equivalent positions.

where $\hat{\mathbf{z}}_j$ is the local $[111]$ axis, which is the axis connecting the centers of the two tetrahedra that share the sites. In this case, the constraint $M_t = 0$ is equivalent to a “two-in-two-out” rule for the moments $\boldsymbol{\mu}$, so that two of the four moments point into each tetrahedron and two point out of it. One such example configuration is shown in Fig. 7.1 (b), alongside a possible magnetic field line configuration indicated in blue. Taking the coarse-grained field \mathbf{B} to be equal to the moment $\boldsymbol{\mu}_j$ on site j , the two-in-two-out rule forces the number of magnetic field lines entering the tetrahedron to be equal to those leaving, corresponding to a Gauss’s law enforcing zero divergence

$$\nabla \cdot \mathbf{B} = 0 \quad (7.4)$$

The ground state manifold of Eq. (7.1) is hence naturally described in terms of an emergent Maxwell gauge theory. Violations the ground-state constraint $M_h = 0$ are violations of the Gauss’s law in Eq. (7.4) and hence allow for a natural interpretation as magnetic monopoles [120]. A single spin flip creates a pair of monopoles, both of which can move freely through the system by further spin flips. This is an example of fractionalization: while the constituent degrees of freedom are magnetic dipole moments $\boldsymbol{\mu}_j$, the emergent excitations take the form of monopoles. The possibility to separate charges at finite energy cost means that the emergent Maxwell theory is in its deconfined phase, the so-called the Coulomb phase.

7.1.2. Spin ice materials

Given the third law of thermodynamics stating that every system must have a well-defined, constant entropy in the limit $T \rightarrow 0$, a natural question to ask is whether the above scenario can be realized in the real world. Indeed, the model in Eq. (7.1) is fine-tuned in the sense that any added generic perturbation will lead to ordering at low temperature. It turns out however, that a family of materials exist, in which the above mentioned physics is realized in a quite well-defined regime at

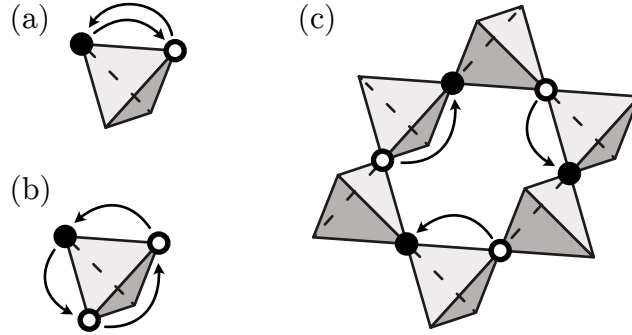


Figure 7.3.: Processes contributing to third-order degenerate perturbation theory for quantum spin ice, that is Eq. (7.6) with $J_{\perp} \ll J_{zz}$. The processes shown in (a) and (b) lead only to constant energy shifts while (c) connects different configurations.

finite temperature [203–205]. In particular, the Gauss’s law in Eq. (7.4) upon Fourier transform gives rise to a characteristic “pinch-point” singularity of the magnetic structure factor [116, 121, 122, 206, 207]

$$S(\mathbf{q}) = \sum_{\alpha,\beta} \left(\delta_{\alpha\beta} - \frac{q^{\alpha} q^{\beta}}{q^2} \right) \sum_{i,j} e^{-i(\mathbf{r}_i - \mathbf{r}_j) \cdot \mathbf{q}} \langle \mu_i^{\alpha} \cdot \mu_j^{\beta} \rangle, \quad (7.5)$$

where \mathbf{r}_j is the position of the magnetic ion j . This quantity is shown in Fig. 7.2, with pinch-points clearly visible at the zone centers $\mathbf{q} = [1, 1, 1]$. Importantly, this characteristic pattern has also been observed in neutron scattering experiments on $\text{Ho}_2\text{Ti}_2\text{O}_7$ [206] and $\text{Dy}_2\text{Ti}_2\text{O}_7$ [207], confirming the realization of the Coulomb phase in this material.

Spin ice materials provide a fruitful avenue to discover new physics, for theorists and experimentalist alike, because of reliable experimental data available together with a detailed understanding of which theory models are relevant for capturing both, equilibrium properties as well as — more recently — their dynamics [208].

Here, we will not go into much detail describing all riches available in the field but instead use the description above as a motivation to study the quantum version of the problem, called quantum spin ice. As we will discuss in the following, the existence of quantum spin ice is well established in theory, but it has not been conclusively observed in any material setting.

7.2. Quantum Spin Ice

Classical spin liquids are intrinsically unstable to perturbations because of their residual entropy, that violates the third law of thermodynamics. While the existence of such a residual entropy may be used as a good definition of such classical spin liquids, these systems have many other interesting properties such as emergent gauge theories and fractionalized excitations. In a classical system the existence of a local symmetry already implies an extensive entropy. This is not the case in the quantum setting, implying the possibility to stabilize a quantum analog of classical spin liquids by means of quantum fluctuations. As sketched in essence already in Anderson’s 1973 proposal of the “resonating valence bond state” [12], the idea

is that fluctuations between classically degenerate ground states stabilize a macroscopic superposition of these states as the ground state of the quantum system.

7.2.1. A quantum liquid

As a minimal model, consider for example the XXZ model on the pyrochlore lattice

$$H = \underbrace{J_{zz} \sum_{\langle ij \rangle} \hat{S}_i^z \hat{S}_j^z}_{H_{zz}} - \underbrace{J_{\perp} \sum_{\langle ij \rangle} (\hat{S}_i^+ \hat{S}_j^- + \hat{S}_i^- \hat{S}_j^+)}_{H_{\perp}} \quad (7.6)$$

where H_{zz} is the classical spin ice Hamiltonian discussed before, and for small transverse exchange $J_{\perp} \ll J_{zz}$, the off-diagonal part H_{\perp} can be treated within degenerate perturbation theory [196]. Different classical spin configurations are only connected at third-order. Denoting by \mathcal{P}_0 the projector onto the ground-state manifold, we find an effective Hamiltonian

$$H_{\text{eff}} = \mathcal{P}_0 \left[-H_{\perp} \frac{1 - \mathcal{P}_0}{H_{zz}} H_{\perp} + H_{\perp} \frac{1 - \mathcal{P}_0}{H_{zz}} H_{\perp} \frac{1 - \mathcal{P}_0}{H_{zz}} H_{\perp} \right] \mathcal{P}_0 \quad (7.7)$$

$$= \frac{J_{\perp}^2}{J_{zz}} \left(\frac{J_{\perp}}{J_{zz}} - 1 \right) N_t + J_{\square} \sum_{\square} \left(\hat{S}_1^+ \hat{S}_2^- \hat{S}_3^+ \hat{S}_4^- \hat{S}_5^+ \hat{S}_6^- + \text{h.c.} \right), \quad (7.8)$$

with $J_{\square} = (3J_{\perp}^3)/(2J_z^2)$ and the sum is over all hexagons \square in the lattice. For each hexagon, there is a ring exchange term with the $\hat{S}_1^+ \hat{S}_2^- \hat{S}_3^+ \hat{S}_4^- \hat{S}_5^+ \hat{S}_6^-$ denoting a product along the hexagon. The first, constant term in the second line is generated by the processes shown in Fig. 7.3 (a) and (b), while the ring exchange term is generated by the resonance shown in panel (c) of the same figure.

It is not immediately clear what is the ground state of the effective Hamiltonian in Eq. (7.8). However it is intimately connected to the physics of quantum dimer models as studied by Rokhsar and Kivelson [194] and can be tuned to an exactly solvable point in a similar manner. In particular, consider adding a potential term of the form

$$H_{\text{flip}} = V N_{\text{flip}} = V \sum_{\square} \mathcal{P}_{\text{flip}} \quad (7.9)$$

which counts the number of “flippable” hexagons N_{flip} on which the ring exchange term acts non-trivially. In this case, we can write

$$H_{\text{eff}} + H_{\text{flip}} = \sum_{\square} J_{\square} \left(\hat{S}_1^+ \hat{S}_2^- \hat{S}_3^+ \hat{S}_4^- \hat{S}_5^+ \hat{S}_6^- + \text{h.c.} \right) + V \mathcal{P}_{\text{flip}} \quad (7.10)$$

$$= \sum_{\square} J_{\square} (|\bullet\bullet\bullet\rangle\langle\bullet\bullet\bullet| + |\bullet\bullet\bullet\rangle\langle\bullet\bullet\bullet|) + V (|\bullet\bullet\bullet\rangle\langle\bullet\bullet\bullet| + |\bullet\bullet\bullet\rangle\langle\bullet\bullet\bullet|) \quad (7.11)$$

where we have dropped the constant terms in Eq. (7.8). For $V = J_{\square}$, the ground state is given exactly as the equal-amplitude superposition of a set of all classical spin ice states connected by possible applications of the ring exchange term

$$|RK\rangle = \prod_{\square} (1 + |\bullet\bullet\bullet\rangle\langle\bullet\bullet\bullet| + |\bullet\bullet\bullet\rangle\langle\bullet\bullet\bullet|) |\Psi_0\rangle, \quad (7.12)$$

for a given classical spin ice configuration $|\Psi_0\rangle$. To see this, note that in this case we can write

$$H_{\text{eff}} + H_{\text{flip}} = V \sum_{\square} (|\langle \bullet \bullet \bullet \rangle - |\langle \bullet \bullet \bullet \rangle\rangle) (\langle \langle \bullet \bullet \bullet | - \langle \langle \bullet \bullet \bullet |), \quad (7.13)$$

which is a sum of projectors with positive coefficient. It also annihilates the state $|RK\rangle$, making this an eigenstate with eigenvalue zero. Assuming ergodicity of the ring exchange term within the set of classical spin ice states, the argument above indeed established the existence of a *quantum* spin liquid state at least at the exactly solvable point (also called the Rokhsar-Kivelson point). As conjectured in the beginning of the section, its wavefunction takes the form of a superposition of all classical spin ice states within a given sector connected by the ring exchange term. An open question at this point is the extent of this quantum spin ice phase beyond the exactly solvable point. We will return to this question after deriving an effective gauge theory describing the phase in the following.

7.2.2. Emergent $U(1)$ gauge theory

The quantum spin ice state derived above has an emergent gauge theory structure at low energy. To see this, we follow the discussion in Ref. [196] and rewrite the effective Hamiltonian in Eq. (7.8) as a pure gauge theory on the lattice. In particular, we introduce quantum rotor variables

$$\hat{n}_j \in \mathbb{Z}, \quad \hat{\phi} \in [-\pi, \pi), \quad [\hat{\phi}_i, \hat{n}_j] = i\delta_{ij}. \quad (7.14)$$

in term of which $\hat{S}^z = \hat{n} - \frac{1}{2}$ and $\hat{S}^\pm = \exp(\pm i\hat{\phi})$. Hence, we can rewrite the effective Hamiltonian (omitting constant terms) as

$$H_{\text{eff}} = U \sum_j \left(\hat{n}_j - \frac{1}{2} \right)^2 + 2J_{\square} \sum_{\square} \cos(\hat{\phi}_1 - \hat{\phi}_2 + \hat{\phi}_3 - \hat{\phi}_4 + \hat{\phi}_5 - \hat{\phi}_6) \quad (7.15)$$

in the limit $U \rightarrow \infty$, which enforces the hardcore constraint $\hat{n} \in \{0, 1\}$ for spin-1/2. Similar to the case of classical spin ice, we can now define an electric field \mathbf{E} and a vector potential \mathbf{A}

$$\mathbf{E}(\mathbf{r}_j) = \mathbf{z}_j \left(\hat{n}_j - \frac{1}{2} \right), \quad \mathbf{A}(\mathbf{r}_j) = \mathbf{z}_j \hat{\phi}_j \quad (7.16)$$

such that

$$H_{\text{eff}} = U \sum_j \mathbf{E}(\mathbf{r}_j)^2 + 2J_{\square} \sum_{\square} \cos(\text{curl}_{\square} \mathbf{A}) \quad (7.17)$$

where

$$\text{curl}_{\square} \mathbf{A} = \int_{\square} \mathbf{A} \cdot d\mathbf{l} \quad (7.18)$$

can be thought of as the discrete line integral of \mathbf{A} along the smallest closed loop in the lattice, a hexagon. Using Stokes' theorem, this is equal to the magnetic flux Φ_{\square} through the hexagon \square .

The Hamiltonian in Eq. (7.17) is that of *compact* quantum electrodynamics in three dimensions (c-QED₃). Note that in the usual notation of QED₃, what we




Figure 7.4.: Phase diagram for (a) the RK-Hamiltonian in Eq. (7.11) and (b) the XXZ model in Eq. (7.6). The point $V = 0$ in the RK model corresponds to taking $J_{zz} \rightarrow \infty$ at finite J_{\perp} .

called the magnetic field in classical spin ice, is now the electric field of quantum spin ice [209]. This can lead to some confusing notation, given that the electric charge of the theory

$$(\nabla \cdot \mathbf{E})_{\text{tet}} = \begin{cases} \sum_{j \in \text{tet}} S_j^z & \text{up tetrahedra} \\ -\sum_{j \in \text{tet}} S_j^z & \text{down tetrahedra} \end{cases}, \quad (7.19)$$

corresponds to the magnetic monopoles that we described in the section on classical spin ice. In particular, since the value of S^z in real materials corresponds directly to the physical magnetic moment [cf. Eq. (7.3)], this “electric charge” carries an actual magnetic charge [120]. Since the emergent QED theory is compact, it has another kind of gapped topological excitation, which is in fact usually called magnetic monopole, as it corresponds to vortices in the magnetic field \mathbf{B} . To avoid confusion we will adopt a slightly different notation throughout this thesis, and call violations of the spin ice rule (the electric charges in c-QED₃) magnetic monopoles, and call the vortices of the emergent electric field visons.

For completeness, note that there is a third type of excitation, that is a gauge boson called the photon. Since the gauge symmetry is continuous, it is gapless and has been proposed as a characteristic signature of the quantum spin ice phase since it is in principle observable in inelastic neutron scattering [210]. Because the photon is gapless, it also leads to a characteristic low-temperature specific heat $C_v \sim (T/v_{\text{ph}})^3$, where v_{ph} is the photon velocity. For the same reason, its presence washes out the pinch point singularity in the structure factor [210], which was introduced as a characteristic signature of classical spin ice in Sec. 7.1 [cf. Fig. 7.2].

Note that the Hamiltonian in Eq. (7.17) admits two distinct deconfined phases, depending on the sign of the ring exchange $J_{\square} \propto J_{\perp}^3/J_z^2$. Since it is the prefactor of the cosine term, the magnetic sector in the ground state either fulfills $\Phi_{\square} = 0$ if $J_{\square} < 0$ or $\Phi_{\square} = \pi$ if $J_{\square} > 0$. These two scenarios are called zero and π flux, respectively. The latter, also denoted as $U(1)_{\pi}$ phase, is characterized by the fact that when a monopole hops around a single hexagon, the total wave function picks up a minus sign. As a consequence, the monopole dispersion has enhanced periodicity compared to the zero-flux $U(1)_0$ phase, which is observable in inelastic neutron scattering [211].

7.2.3. Phase diagram of the XXZ-model

Compact quantum electrodynamics in three dimension at zero temperature is stable even to perturbations breaking the emergent, continuous $U(1)$ gauge invariance

[196, 209]. This is in contrast to the situation in two dimensions, where any compact $U(1)$ gauge theory with gapped matter confines [181]. Considering the Hamiltonian $H_{\text{eff}} + H_{\text{flip}}$ in Eq. (7.11) on the pyrochlore lattice, one can hence expect that the quantum spin ice phase has a finite extent beyond the exactly solvable point $V = J_{\square}$. Whether the phase extends all the way to $V = 0$, which corresponds to the XXZ spin model in Eq. (7.6) is not clear. Hence, the existence of a low-temperature QSI regime in this model can only be established numerically. Fortunately, at least for ferromagnetic transverse exchange $J_{\perp} < 0$, the XXZ Hamiltonian does not have a sign problem (that is all its off-diagonal matrix elements are negative) and allows for exact numerical solution using quantum Monte-Carlo (QMC). Due to such simulations, the existence of a zero-temperature QSI phase in the XXZ model is well established for $J_{\perp}/J_z \lesssim 0.05$ [212–215]. For larger values of J_{\perp} , the system orders into a transverse XY-ferromagnet.

Note that the cases $J_{\perp} < 0$ and $J_{\perp} > 0$ are equivalent in third order perturbation theory, that is, there is unitary transformation taking $J_{\square} \rightarrow -J_{\square}$ in Eq. (7.8) [196]. Because of this, we expect also the $U(1)_{\pi}$ liquid to be stable for some finite parameter region of antiferromagnetic transverse exchange $J_{\perp} > 0$.

These results are also summarized in Fig. 7.4. Panel (a) shows a sketch of the phase diagram of Eq. (7.11). The QSL phase extends all the way from the exactly solvable RK point to $V = 0$, which corresponds to the $J_{zz} \rightarrow \infty$ limit of the XXZ spin model. The phase diagram of this model for finite J_{zz} is shown in panel (b) of the same figure. Around the classical spin ice point $J_{\perp} = 0$, there are two $U(1)$ -QSL phases with zero and π flux for $J_{\perp} < 0$ and $J_{\perp} > 0$, respectively. While the extent of the $U(1)_0$ phase is known exactly from QMC simulations, the fate of the $U(1)_{\pi}$ phase for non-perturbative $J_{\perp} > 0$ remains an open question.

7.3. Dipolar-Octupolar Pyrochlore Materials

In the previous section, we described how quantum spin ice (QSI) emerges in the low-temperature regime of a XXZ model on the pyrochlore lattice. While this established the existence of the phase in this particular model, it is not clear that this model is realized in any real material. As it turns out, there is a large class of materials called dipolar-octupolar pyrochlores [216, 217], which are at low temperature well described by a slight generalization of the XXZ model on the pyrochlore lattice. In the following, we first sketch the derivation of the most general nearest-neighbor exchange Hamiltonian realized in this class of materials. In the second part of the section, we discuss the phase diagram of this model and its current theoretical understanding.

7.3.1. Hierarchy of energy scales

The pyrochlore lattice is quite ubiquitous in nature and a large number of magnetic compounds realizing it have been synthesized. A large class of them, so called rare-earth pyrochlores, follow the chemical formula $R_2M_2O_7$, where R is a rare-earth in the 3+ oxidation state, M is a transition metal in the 4+ state and O is oxygen. Compounds in this family include, among many others, the classical spin ices $\text{Ho}_2\text{Ti}_2\text{O}_7$ and $\text{Dy}_2\text{Ti}_2\text{O}_7$, the material $\text{Er}_2\text{Ti}_2\text{O}_7$ mentioned in Sec. 2.1.2 as a well established example of quantum order-by-disorder, as well as the Cerium-based




Figure 7.5.: Classification of crystal field doublets. The site symmetry group of the rare-earth ions R^{3+} , D_{3d} , is generated by one three-fold rotation C_3 shown in (a), three distinct mirror symmetries σ shown in (b), together with inversion (not shown). As shown in (c), all possible ground state doublets of the crystal field can be classified by their transformation properties under D_{3d} together with time reversal T .

quantum spin ice candidates $\text{Ce}_2\text{M}_2\text{O}_7$ ($M=\text{Zr},\text{Sn},\text{Hf}$) that will be the subject of this chapter.

In all cases, studying the material in theory requires establishing an effective low-energy exchange Hamiltonian. This is possible since rare-earth magnets feature a reasonably clear hierarchy of energy scales [217], with single-ion physics dominating over two-ion exchange interactions. For the single-ion physics in turn Coulomb interactions and spin-orbit coupling, as summarized by Hund's rules, dominate over crystal field effects.

Taking Cerium as an example as it is the ion most relevant for this thesis, its 58 electrons are arranged in the configuration $[\text{Xe}]4f^15d^26s^2$. The trivalent oxidation stage relevant for rare-earth pyrochlores then has a single 4f valence electron, with total spin $S = 1/2$ and orbital angular momentum $L = 3$. Since spin-orbit coupling is comparatively strong in rare-earth ions, the good quantum number, however, is total angular momentum, which according to Hund's rules takes the value $J = L - S = 5/2$.

7.3.2. Classification of crystal field doublets

Most of the remaining degeneracy of the levels $|J = \frac{5}{2}, m_J = -\frac{5}{2} \dots \frac{5}{2}\rangle$ is then lifted by the crystal electric field, that is by the effect of the crystalline environment. In all cases considered here, the ground state of the crystal field Hamiltonian is a doublet, which we denote by $|\pm\rangle$. Then, one can define pseudospin- $\frac{1}{2}$ operators as

$$\hat{S}^z = \frac{1}{2} (|+\rangle\langle+| - |-\rangle\langle-|), \quad \hat{S}^\pm = |\pm\rangle\langle\mp|. \quad (7.20)$$

Since the crystal field has to fulfill the on-site symmetry, the possible ground state doublets can be classified by their transformation properties under that symmetry. Precisely, they are given by the irreducible representation of the (quantum doublet of) the point-group symmetry plus time reversal. The on-site symmetry group of the pyrochlore lattice D_{3d} is illustrated in Fig. 7.5. It is generated by a C_3 shown in Fig. 7.5 (a), three mirror symmetries shown in panel (b) of the same figure, together with inversion, which is not shown. The C_3 rotation is about the local $[111]$ easy

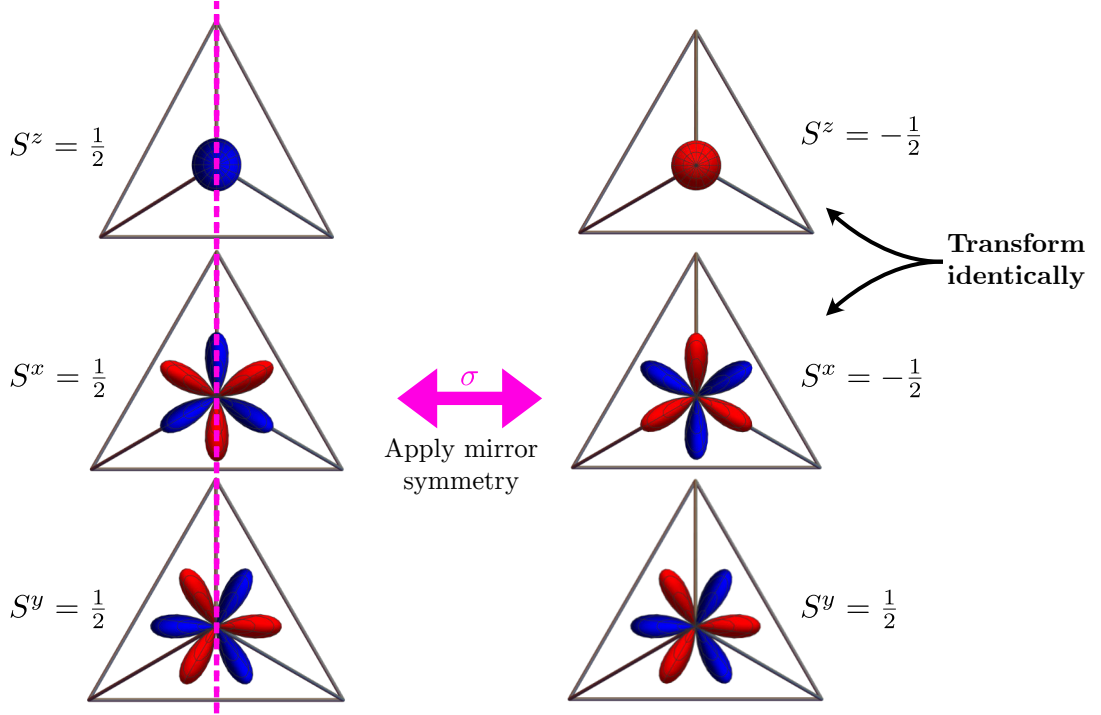


Figure 7.6.: The effect of the local mirror symmetry σ on the magnetic pseudocharge distribution of the eigenstates of the pseudospin components \hat{S}^α .

axis, and the three mirror symmetries are each about a plane cutting through one of the incident bonds as well as through the position of the neighboring transition metal M^{4+} , located at the center of the tetrahedra. For clarity, we do not show the position of the oxygen atoms in Fig. 7.5.

The classification of possible doublets in three classes is then shown in Fig. 7.5 (c). First, Kramers ions with an odd number of electrons and non-Kramers ions with an even number of electrons (that is half-integer and integer J , respectively) are distinguished by their transformation properties under time reversal. Kramers-ions transform non-trivially under time reversal, which protects the two-fold degeneracy, while doublets for non-Kramers ions are protected by lattice symmetry. Kramers ions can be further classified into two classes, distinguished by how they transform under spacial symmetries. The first is the effective spin- $\frac{1}{2}$ doublet, which transforms as a local magnetic dipole under all on-site symmetries. The second class, which includes all known Cerium-based compounds, is called dipolar-octupolar doublet since its pseudospin \hat{S}^z operator corresponds to a component local dipole moment while the \hat{S}^\pm components correspond to parts of a magnetic octupole moment [216]. In the known Cerium-based pyrochlores for example, the ground state doublet to a good approximation is given by [218–225]

$$|\pm\rangle_{\text{Ce}^{3+}} = \left| J = \frac{5}{2}, m_J = \pm \frac{3}{2} \right\rangle, \quad (7.21)$$

possibly including a small component with larger total angular momentum $J = \frac{7}{2}$ [219, 224, 226].

In this case, while \hat{S}^z is naturally a dipole, the off-diagonal operators as defined in Eq. (7.20) change the angular momentum in units of three and hence are compo-

nents of an octupole moment. Importantly, dipolar and octupolar moments are not necessarily distinguished by the elements of the on-site symmetry group, as shown in Fig. 7.6. The figure illustrates that while the operator \hat{S}^y transforms trivially under the mirror symmetry σ , the operators \hat{S}^z and \hat{S}^x both transform non-trivially. Indeed, it is easy to show that \hat{S}^z and \hat{S}^x transform identically under all on-site symmetries. Hence, both transverse pseudospin components correspond physically to components of a magnetic octupole moment, but the component S^x transforms like the component of a dipole operator under on-site symmetries.

The fact that parts of the local moment correspond to the component of a magnetic octupole naturally has consequences for the coupling to an external magnetic field \mathbf{H} . In leading order, it couples only to the local dipole moment, given by S^z [226, 227]

$$H_{\text{field}} = \sum_j \boldsymbol{\mu}_j \cdot \mathbf{H}, \quad (7.22)$$

$$\boldsymbol{\mu}_j = -\mu_B \hat{\mathbf{z}}_j g_z \hat{S}_j^z \quad (7.23)$$

where $\hat{\mathbf{z}}_j$ is pointing along the local [111] easy-axis (which is identical to the C_3 rotation axis in Fig. 7.5). The constant μ_B is the Bohr magneton and g_z is the material-dependent g-factor of the ion. For Cerium, when assuming a pure $J = \frac{5}{2}$ crystal field ground state as in Eq. (7.21), $g_z = 18/7 \approx 2.57$

Note that, in principle, also the pseudospin S^x operator is allowed to couple to the external field purely on grounds of symmetry [226, 227]. However, the corresponding g-factor g_x is suppressed by selection rules fulfilled by the matrix elements of the magnetization density. Explicitly, by virtue of the the Wigner-Eckhart theorem for a rank-1 spherical tensor, these matrix elements μ_q , $q = -1, 0, 1$ fulfill

$$\langle J, m | \mu_q | J', m' \rangle \propto \langle J' m'; 1q | Jm \rangle, \quad (7.24)$$

where $\langle J' m'; kq | Jm \rangle$ denotes the Glebsch-Gordan coefficient, which is nonzero only if $|J' - 1| < J < J' + 1$ and $m = m' + q$. Since $m - m' \in 3\mathbb{Z}$ for dipolar octupolar doublets, the only nonzero component of $\boldsymbol{\mu}$ is $q = 0$ (it always points along the local [111] easy axis). Assuming the ideal Cerium doublet in Eq. (7.21), μ_0 is clearly diagonal in $|\pm\rangle$ and hence $\mu_0 \propto S^z$. In general, however, the doublet can include a component with higher angular momentum J , and it is not always possible to choose a basis of the doublet such that μ_0 is diagonal. We can however always choose a basis such that μ_0 is real and hence it is always a combination of only S^z and S^x . As mentioned already above, the mixing of the ideal doublet with states of larger angular momentum, and therefore also the value of g_x , is found to be very small [219, 220].

7.3.3. Exchange Hamiltonian and coupling to external field

The different transformation properties of the magnetic pseudospin components put important constraints on the most general, symmetry allowed nearest-neighbor exchange Hamiltonian. In particular, the only allowed couplings are between components that transform identically

$$H_{\text{ex}} = \sum_{\langle ij \rangle} J_x S_i^x S_j^x + J_y S_i^y S_j^y + J_z S_i^z S_j^z + J_{xz} (S_i^x S_j^z + S_i^z S_j^x) \quad (7.25)$$

where the sum is over the nearest-neighbor bonds $\langle ij \rangle$ of the pyrochlore lattice. For simplicity, in the following we omit the hats on top of operators and instead use hats only to denote unit vectors.

The Hamiltonian in Eq. (7.25) can be further simplified. The XZ exchange can be removed by a rotation of the pseudospin axes by an angle θ around the pseudospin y -axis

$$S_i^x \rightarrow \cos \theta S_i^x - \sin \theta S_i^z, \quad (7.26a)$$

$$S_i^z \rightarrow \sin \theta S_i^x + \cos \theta S_i^z, \quad (7.26b)$$

$$\tan(2\theta) = \frac{2J_{xz}}{J_x - J_z}, \quad (7.26c)$$

which yields

$$H = H_{\text{ex}} + H_{\text{field}}, \quad (7.27a)$$

$$H_{\text{ex}} = \sum_{\langle ij \rangle} J_x S_i^x S_j^x + J_y S_i^y S_j^y + J_z S_i^z S_j^z, \quad (7.27b)$$

$$H_{\text{field}} = -\mu_B g_z \sum_j \mathbf{H} \cdot \hat{\mathbf{z}}_j (\sin \theta S_j^x + \cos \theta S_j^z). \quad (7.27c)$$

To avoid too ornamented notation, we do not rename the spin operators S_i^α , but instead from here on the S_i^α refer to those in Eq. (7.27). Note that this transformation is also possible if starting from a Hamiltonian with nonzero value of g_x . In the parametrization used above, this leads to a renormalization of the value of g_z compared to the value in the pure $J = 5/2$ doublet.

Altogether, dipolar-octupolar pyrochlores are described by a relatively simple XYZ exchange Hamiltonian with a peculiar coupling to the magnetic field. Expanding on the discussion of the XXZ model in the previous section, in the following we give a brief overview of the current understanding of the phase diagram of the XYZ model.

7.3.4. Phase diagram of dipolar-octupolar pyrochlores

For now, we are interested in the phase diagram of Eq. (7.27) without a magnetic field, $\mathbf{H} = 0$. To streamline the discussion, it is useful to introduce a change of variables naturally accounting for the symmetries of H_{ex} . Consider a permutation of the couplings (J_x, J_y, J_z) , which we call (J_a, J_b, J_c) such that $|J_a| > |J_b|, |J_c|$ and $J_b > J_c$. To obtain the full phase diagram as a function of (J_x, J_y, J_z) , it is hence sufficient to study the phase diagram as a function of the symmetry-reduced parameters (J_a, J_b, J_c) . We also rewrite the Hamiltonian in terms of ladder operators

$$H = \sum_{ij} J_a S_i^a S_j^a - J_\pm (S_i^+ S_j^- + S_i^- S_j^+) + J_{\pm\pm} (S_i^+ S_j^+ + S_i^- S_j^-), \quad (7.28)$$

where $J_\pm = -\frac{1}{4}(J_b + J_c)$ and $J_{\pm\pm} = \frac{1}{4}(J_b - J_c)$. The conditions on the (J_a, J_b, J_c) then defines a triangle in the parameter space of $(J_a, J_\pm, J_{\pm\pm})$.

The phase diagram of the model in this subspace of parameter space is shown in Fig. 7.7. Note that for $J_{\pm\pm} = 0$, the model in Eq. (7.28) reduces to the XXZ model (with $J_\perp = -J_\pm$) already discussed in Sec. 7.2. There, we already introduced

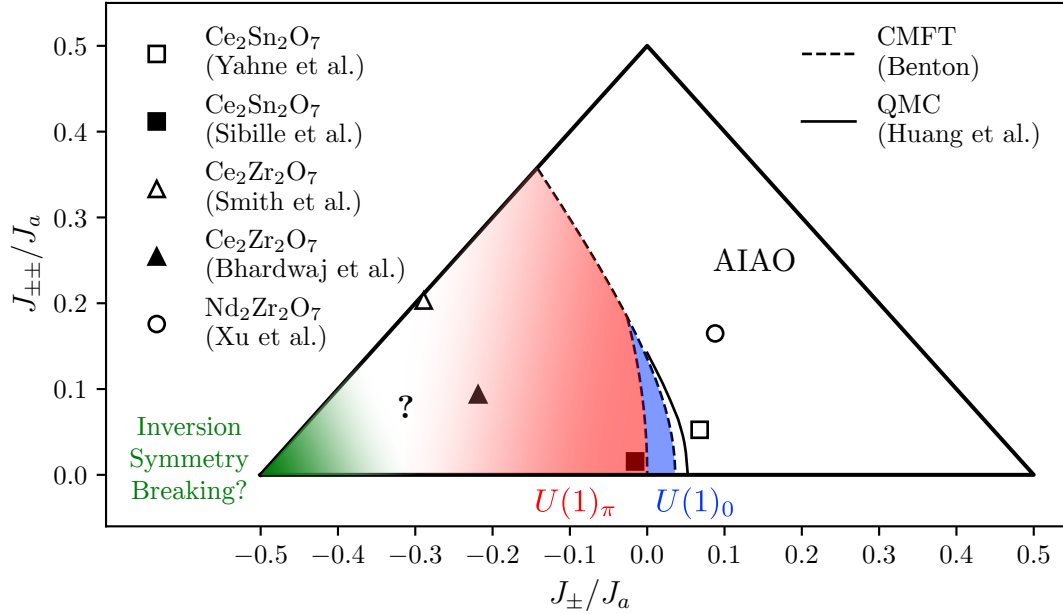


Figure 7.7.: Ground state phase diagram of the XYZ model on the pyrochlore lattice [Eq. (7.28)] as realized in dipolar-octupolar pyrochlores. Three phases are found within cluster mean field theory [228]: magnetic all-in-all-out (AIAO) order and two $U(1)_{\pi}$ quantum spin liquids with zero and π flux (cf. Sec. 7.2). Dashed and solid lines indicated the phase boundaries as found using CMFT [228] and quantum Monte-Carlo [215]. Since CMFT is an uncontrolled method, color gradients suggest the degree of certainty that a given ground state is realized. Strong numerical evidence for inversion symmetry breaking at the Heisenberg point was found recently using DMRG [229]. Markers indicate various estimates of exchange parameters for dipolar-octupolar pyrochlore materials [92, 219, 222, 226, 230].

most of the ground states appearing in Fig. 7.7. The point $J_{\pm} = J_{\pm\pm} = 0$ is just classical spin ice. Introducing quantum fluctuations as discussed in Sec. 7.2 drives the system into one of two possible $U(1)$ quantum spin liquid ground states, with zero and π flux around the hexagon for $J_{\pm} > 0$ and $J_{\pm} < 0$, respectively. For large ferromagnetic transverse exchange, the system develops transverse ferromagnetic long range order. Since the local magnetic moment is always parallel to the local easy axis, this corresponds to an “all-in-all-out” state where for each tetrahedron all magnetic moments either point into the tetrahedron or out of it.

For unfrustrated (that is ferromagnetic) transverse exchange, $J_{\pm} > 0$, the Hamiltonian in Eq. (7.28) still allows for exact numerical simulation using quantum Monte-Carlo. Because of this, the phase boundary between the all-in-all-out phase and the zero-flux quantum spin ice state is known exactly [215], and indicated in Fig. 7.7 by a solid line. The authors of Ref. [215], also claim the existence of a \mathbb{Z}_2 QSL phase, based on the scaling of the low-temperature specific heat. Concretely, they observe an exponential decay of the specific heat at low temperatures, in a narrow region of parameter space sandwiched between the $U(1)_0$ and the AIAO phase and for small but nonzero $J_{\pm\pm} > 0$. An exponentially suppressed specific heat at low temperature is indeed inconsistent with the existence of an emergent photon excitation expected in a $U(1)_0$ quantum spin ice state, but further studies are needed to verify the existence of a \mathbb{Z}_2 QSL in this region.

For frustrated transverse exchange, the only available insight into the phase diagram are based on mean-field calculations. Perhaps the most reliable is the cluster mean field (CMFT) study in Ref. [228]. The phase boundary derived there is shown as a dashed line in Fig. 7.7 and agrees reasonably well with the QMC result. However, CMFT still constitutes an uncontrolled approximation and hence the derived phase diagram comes with a significant uncertainty, especially in the part of the phase diagram with lacks long-range magnetic order. Because of this, in Fig. 7.7 color gradients are used as a rough guide towards the degree of certainty of the phase diagram as drawn.

CMFT predicts an extended $U(1)_{\pi}$ spin liquid state all the way to the Heisenberg point $J_{\pm} = -\frac{1}{2}$. The π -flux phase is indeed expected to be more stable than the 0-flux phase because in the former case also the transverse exchange is frustrated. A multitude of proposals exist for the ground state of the spin- $\frac{1}{2}$ Heisenberg model on the pyrochlore lattice, including various forms of quantum spin liquid [231–236], valence-bond solids [229, 237–245], or the possibility that it lies on a phase boundary between different QSLs [246]. While its nature remains to be settled conclusively, the most reliable results available are probably the large-scale DMRG studies of Refs. [229, 242] with the most robust result being spontaneously broken inversion symmetry. This is in accordance with a number of follow up studies [243, 244], including the recent suggestion of a particular valence bond solid taking the form of a resonating hexagons [245]. Since a valence bond solid is gapped, one can expect such a state to be stable over a finite region of parameter space. The fact that it is unclear how big the extent of such a non-magnetically ordered phase would be is again suggested by a color gradient in Fig. 7.7.

In summary, while a big part of the phase diagram is understood rather well, there is still a significant uncertainty about the ground state of the spin- $\frac{1}{2}$ model for large frustrated transverse exchange.

Note that while all permutations of (J_x, J_y, J_z) are equivalent in terms of thermo-

dynamic properties, they do differ quite drastically with regards to their magnetic properties. This is because only the operator S^z corresponds to a magnetic dipole, while the other two pseudospin components correspond to components of a magnetic octupole moment. Accordingly, each of phases in Fig. 7.7 has two incarnations, one where the longitudinal pseudospin operator S^a corresponds to the component of a dipole and one where it corresponds to part of an octupole moment [228, 247]. In the case of the AIAO phase, the two incarnations simply translates into the nature of the ordered moment. For the non-magnetic phases it dictates the moment carried by its excitations. In all cases, this has consequences on the signatures of the phases observable in neutron scattering, since the neutron at low momentum transfer only couples to the local dipole moment. This will be discussed in detail in Sec. 8.1. It also impacts the response of the phase in an external magnetic field [247, 248], as will be discussed in chapter 9.

In Fig. 7.7, markers indicate also various estimates of exchange parameter of dipolar-octupolar pyrochlore materials. These are obtained by fitting theory calculations to thermodynamic measurements in the case of $\text{Ce}_2\text{Zr}_2\text{O}_7$ [222, 226], and $\text{Ce}_2\text{Sn}_2\text{O}_7$ [218, 223], and by fitting spin wave dispersions as measured in inelastic neutron scattering experiments for $\text{Ns}_2\text{Zr}_2\text{O}_7$ [230]. As shown in the figure, different estimates do not necessarily agree even for the same material, owing to the fact experiments as well as theory calculation on three dimensional materials in the quantum regime are very challenging. Part of this thesis were model calculations of neutron scattering and specific heat used in the derivation of two of the estimates shown in Fig. 7.7, Yahne *et al.* and Smith *et al.*. A detailed discussion of how these estimates were obtained in the absence of spin waves will be the subject of chapter 8.

As is also clear from Fig. 7.7, different materials can have very different parameters, even if they are chemically quite similar such as $\text{Ce}_2\text{Zr}_2\text{O}_7$ and $\text{Ce}_2\text{Sn}_2\text{O}_7$. This makes dipolar-octupolar pyrochlores a promising family of materials to finally realize quantum spin ice. An important ingredient is the comparatively simple exchange Hamiltonian, taking the form of an XYZ model on the pyrochlore lattice [Eq. (7.27)]. This allows reliable comparison to theory calculations, which is crucial in order to understand the physics of real materials and to identify signatures of the quantum spin ice phase.

8. Ground State Selection in Dipolar Octupolar Pyrochlores at Zero field

Parts of the work presented in this chapter has appeared in the following articles:

[222] E. M. Smith, O. Benton, D. R. Yahne, B. Placke, et al. “Case for a $U(1)_\pi$ Quantum Spin Liquid Ground State in the Dipole-Octupole Pyrochlore $Ce_2Zr_2O_7$ ”. In: *Phys. Rev. X* 12 (2 Apr. 2022), p. 021015. DOI: 10.1103/PhysRevX.12.021015

[223] D. R. Yahne, B. Placke, et al. *Dipolar spin ice regime proximate to an all-in-all-out Néel ground state in the dipolar-octupolar pyrochlore $Ce_2Sn_2O_7$* . arXiv:2211.15140. 2022. DOI: 10.48550/arXiv.2211.15140

As discussed in the previous chapter, dipolar-octupolar pyrochlores are a promising family of materials to realize quantum spin ice. This is because their most general, symmetry-allowed exchange Hamiltonian takes the relatively simple form of a XYZ model on the pyrochlore lattice. This Hamiltonian is expected to realize quantum spin liquid phases in a large region of its parameter space. Its simple form with only three free parameters also allows for the reliable modeling and in particular the fitting of experimental measurements to theory calculations over the full parameter space.

This modeling, for the materials $Ce_2Zr_2O_7$ and $Ce_2Sn_2O_7$ at zero field, is the subject of this chapter. No sharp spin wave modes have been observed in either of these compounds. Because of this, we extract exchange parameters by fitting various thermodynamic quantities. In the case of $Ce_2Zr_2O_7$, we constrain the exchange Hamiltonian to lie within a region of parameter space expected to realize an exotic $U(1)_\pi$ spin liquid phase. For $Ce_2Sn_2O_7$, a similar fitting procedure together with fits to the powder-averaged structure factor suggests that the ground state of the material realizes all-in-all-out order, but with parameters close to the phase boundary to the $U(1)_0$ spin liquid. Because of this, the critical temperature as predicted by quantum Monte-Carlo simulations is below the regime accessible to current experiments. Experiments hence probe $Ce_2Sn_2O_7$ in a finite-temperature spin ice regime above a long-range ordered state at low-temperature.

The contribution of this thesis to the results described above is to provide theory predictions for both elastic and inelastic neutron scattering measurements, as well as quantum Monte-Carlo predictions for the measured specific heat in the part of the phase diagram where transverse exchange is ferromagnetic.

The rest of the chapter is structured as follows. In Sec. 8.1, we first discuss the neutron scattering response of dipolar-octupolar pyrochlores in general, focussing on

how to incorporate effects beyond the dipole approximation, and the suppression of spin waves even in the field-polarized state for certain model parameters. Building on this, in Sec. 8.2 we describe how to extract model parameters nevertheless by fitting thermodynamic measurements together with neutron scattering, using the material $\text{Ce}_2\text{Zr}_2\text{O}_7$ as an example. In Sec. 8.3, a similar procedure is used to place $\text{Ce}_2\text{Sn}_2\text{O}_7$ into the phase diagram. In this case, the best-fitting exchange parameters admit simulation using quantum Monte Carlo, which enables an even more detailed comparison between theory and experiment. Finally, Sec. 8.4 contains the conclusion and outlook towards some open questions and future work.

8.1. Scattering Neutrons off Dipolar-Octupolar Pyrochlores

Neutron scattering has been proven invaluable in the experimental study of magnetic materials because it is a direct probe of magnetic correlations within the sample. The magnetic contribution to the neutron scattering cross section is given by [249]

$$\left(\frac{d\sigma}{d\Omega}\right)_{\text{mag}} = (\gamma r_0)^2 S(\mathbf{q}) \quad (8.1)$$

where \mathbf{q} is the momentum transferred from the neutron to the sample, and $S(\mathbf{q})$ is the magnetic structure factor

$$S(\mathbf{q}) = \sum_{\alpha,\beta} \left(\delta_{\alpha\beta} - \frac{q^\alpha q^\beta}{q^2} \right) \langle \overline{M^\alpha(\mathbf{q})} \cdot M^\beta(\mathbf{q}) \rangle, \quad (8.2)$$

where $\bar{\bullet}$ denotes complex conjugation and $\mathbf{M}(\mathbf{q})$ denotes the Fourier transform of the local magnetization operator. The magnetic scattering length γr_0 is given by

$$\gamma r_0 = \frac{m_N g_N \mu_N \mu_B \mu_0}{2\pi \hbar^2} \approx 5.39 \times 10^{-15} \text{ m}, \quad (8.3)$$

with m_N the Neutron mass, g_N the g-factor of the Neutron, μ_N the nuclear magneton, μ_B the Bohr magneton, μ_0 the vacuum permeability, and \hbar the reduced Planck constant.

Within the usual dipole approximation, the magnetization operator for dipolar-octupolar pyrochlores takes the form

$$\mathbf{M}(\mathbf{q}) = f_{\text{dip}}(q) \boldsymbol{\mu}(\mathbf{q}) \quad (8.4)$$

where $\boldsymbol{\mu}(\mathbf{q}) = \sum_j \exp(-i\mathbf{q} \cdot \mathbf{r}_j) \boldsymbol{\mu}_j$ is the Fourier transform of the local magnetic dipole moment $\boldsymbol{\mu}_j$ given in Eq. (7.23), and the dipole form factor is

$$f_{\text{dip}}(q) = \langle j_0(q) \rangle + \frac{2 - g_J}{g_J} \langle j_2(q) \rangle, \quad (8.5)$$

with j_n denoting the spherical Bessel functions and g_J the Landé g-factor.

While the above constitutes a good approximation in many cases, in the case of dipolar-octupolar pyrochlores the magnetic dipole moment $\boldsymbol{\mu}$ is only sensitive to part of the pseudospin degree of freedom. In particular, dipolar correlations can be weak or featureless even though strong correlations are present in the system. As will be explained in the following, the dipole contribution, however, is only the leading term in a series of multipole contributions to the total cross section, with larger multipoles contributing only at large momentum transfer.

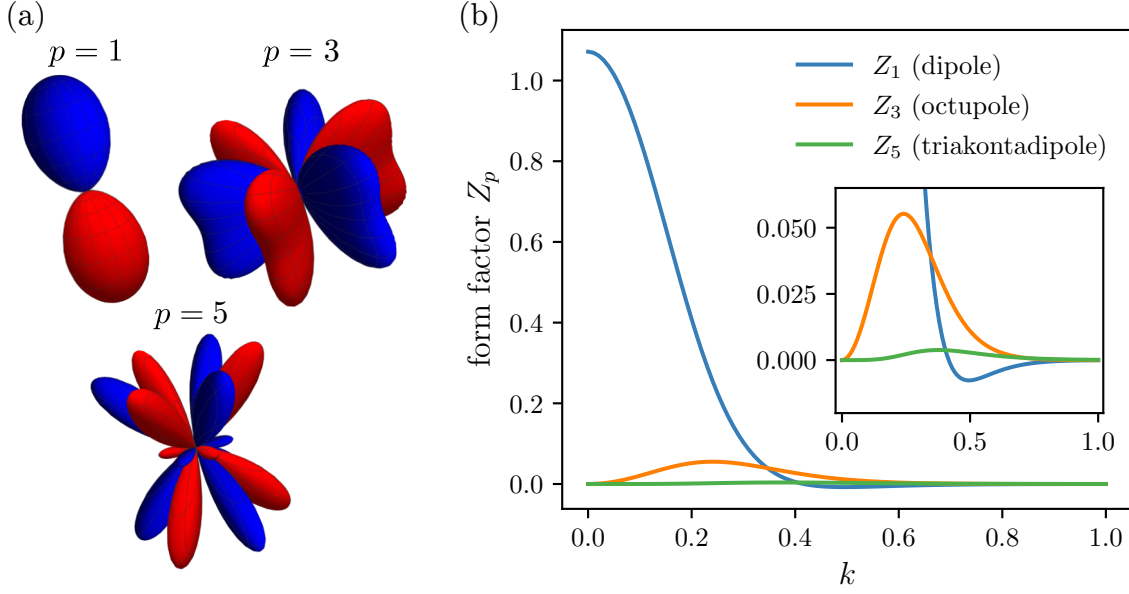


Figure 8.1.: Multipole contribution to the neutron scattering cross section. The magnetization density at each order is the gradient of the pseudocharge distributions shown in (a). The radial dependence is given by the radial integrals Z_p shown in panel (b) for cerium within the hydrogen approximation with $Z_{\text{eff}} = 17$. The reduced momentum is $k = a_0|\mathbf{q}|/Z_{\text{eff}}$ where a_0 is the Bohr radius. Even orders do not contribute due to symmetry.

8.1.1. Neutron scattering cross section beyond the dipole approximation

To express the neutron scattering cross section [Eq. (8.1)] in terms of pseudospin operators S^α , it is necessary to compute the matrix elements of the magnetization operator \mathbf{M} with respect to angular momentum eigenstates. To this end, we treat the wave function of the 4f electrons in Cerium as hydrogen-like with an effective nuclear charge of $Z_{\text{eff}} = 17$ [219, 250]. Within this approximation, one then finds [219, 249–251]

$$\langle SLJM | M^\alpha(\mathbf{k}) | SLJM' \rangle = \sqrt{4\pi} \sum_{p=1}^{2J} \frac{Z_p(k)}{p} \Phi_p^\alpha(\hat{\mathbf{k}}) \quad (8.6)$$

where \mathbf{k} is the reduced momentum $\mathbf{k} = a_0\mathbf{q}/Z_{\text{eff}}$ with a_0 the Bohr radius, and the sum is over multipoles contributions p . Each term can be written as a product of a radial dependence $Z_p(k)$ and an angular dependence $\Phi_p^\alpha(\hat{\mathbf{k}})$. While the full expression for Eq. (8.6) is given in App. D, we here focus on qualitative features of the different orders.

In particular, the magnetization density at each order is the gradient of a suitably defined magnetic pseudocharge distribution. These distributions for the first three contributing orders are shown Fig. 8.1 and correspond to a magnetic dipole ($p = 1$), an octupole ($p = 3$) and an triakontadipole ($p = 5$). Even orders vanish by symmetry. Assuming the pure $J = \frac{5}{2}$ doublet as the crystal field doublet [cf. Eq. (7.21)], these are *all* contributions contributing to the scattering cross section (cf. App. D). The resulting angular dependence $\Phi_p(\hat{\mathbf{k}})$ is uniform at the dipole order ($p = 1$) but highly anisotropic at higher orders.

The radial dependence of the cross section is shown in panel (b) of figure Fig. 8.1. Since $Z_p \propto k^{p-1}$ at small k , the scattering at low momentum is dominated by the dipolar contribution.

Given the expression for the matrix elements in Eq. (8.6), it is easy to express the magnetization operator $\mathbf{M}(\mathbf{k})$ in terms of the Fourier transform of the pseudospin components $S^\alpha(\mathbf{q})$, which are defined in terms of the same angular momentum Eigenstates $|SLJM\rangle$. Generally, one can then relate the correlations of the magnetization operator in Eq. (8.6) to those of the pseudospin operators

$$S(\mathbf{q}) = \sum_{ij} e^{-i\mathbf{q}(\mathbf{r}_i - \mathbf{r}_j)} \sum_{\alpha\beta\gamma\delta} \left(\delta_{\alpha\beta} \frac{q^\alpha q^\beta}{q^2} \right) \mathcal{D}_i^{\alpha\gamma}(\mathbf{q}) \mathcal{D}_j^{\beta\delta}(\mathbf{q}) C_{ij}^{\gamma\delta}(\mathbf{q}) \quad (8.7)$$

where i and j enumerate the four pyrochlore sublattices at positions \mathbf{r}_i and \mathbf{r}_j respectively and

$$C_{ij}^{\gamma\delta}(\mathbf{q}) = \sum_{\mathbf{x}, \mathbf{x}'} e^{-i\mathbf{q}(\mathbf{x} - \mathbf{x}')} \langle S_i^\alpha(\mathbf{x}) S_i^\beta(\mathbf{x}') \rangle \quad (8.8)$$

is the Fourier transform of the two-point correlation function between pseudospin components γ, δ on sublattices i and j . The sum above is over lattice vectors \mathbf{x} . The generalized form factor $\mathcal{D}_j^{\alpha\gamma}(\mathbf{q})$ now takes the form of a 3×3 matrix on each sublattice j and is defined as

$$M_j^\alpha = \sum_{\gamma} \mathcal{D}_j^{\alpha\gamma}(\mathbf{q}) S_j^\gamma(\mathbf{q}). \quad (8.9)$$

The form factor can be written as a sum over the different multipole contributions $\mathcal{D}_j = \mathcal{D}_j^{(1)} + \mathcal{D}_j^{(2)} + \mathcal{D}_j^{(3)}$. The first term reproduces the dipole approximation

$$\mathcal{D}_j^{(1)}(\mathbf{q}) \cdot \mathbf{S}(\mathbf{q}) = f_1(q) \hat{\mathbf{z}}_j S_j^z(\mathbf{q}) \quad (8.10)$$

where $\hat{\mathbf{z}}_j$ is the local [111] easy axis and $S_j^z(\mathbf{q})$ is the Fourier transform of the pseudospin z -component on sublattice j . The expression for the radial form factor $f_1(q)$ differs slightly from the expression within the dipole approximation in Eq. (8.5), but is numerically very close. The full expression for $\mathcal{D}_j(\mathbf{q})$ is given in App. D.

8.1.2. Neutron scattering signatures of dipolar-octupolar spin ices

As a first step towards understanding neutron scattering off dipolar-octupolar pyrochlores more quantitatively, we can consider classical spin ice correlations. This is equivalent to taking the classical limit with $J_\pm = J_{\pm\pm} = 0$ and $J_a > 0$ in Eq. (7.28) and in this case, correlations are easily computed using classical Monte-Carlo simulations and plugged into Eq. (8.7). Assuming also $\theta = 0$ for simplicity, there are then three different states SI_α , $\alpha = x, y, z$ corresponding to classical spin ice correlations realized in the pseudospin α component.

The full neutron scattering cross section from each of these three phases, using the form factor as computed for the pure $J = \frac{5}{2}$ cerium doublet, is shown in Fig. 8.2. For $\alpha = z$, it reproduces the familiar result of dipolar spin ice (cf. Fig. 7.2), although plotted over a much larger window of momentum transfer. The cross section of the




Figure 8.2.: Neutron scattering signatures of three different classical spin ice phases, with classical spin ice correlations in the pseudospin- α component. At large momentum transfer, all three possibilities can be distinguished, without the necessity for large angular resolution. The ice rule, in terms of the respective pseudocharge distribution on a single tetrahedron, is illustrated for each case.

two octupolar spin ice phases $SI_{x,y}$ is less familiar. It is much weaker compared to the dipolar case and scattering intensity is shifted towards large momentum transfer.

This shift in scattering cross section to large momentum transfer can already be used as a diagnostic of dominantly dipolar correlations, since an uncorrelated phase is dominated by the dipolar contribution to neutron scattering. Indeed, in [219] the authors claim the observation of an octupolar spin ice state in $Ce_2Sn_2O_7$, purely based on neutron scattering off a powder sample. Such scattering naturally measures the structure factor averaged over momentum directions, but resolves the magnitude of momentum transfer. The authors observe a broad scattering cross section peaked at a momentum transfer of about 8 \AA^{-1} , corresponding to $|\mathbf{q}| \approx 13$ in the units of Fig. 8.2. If verified, this is a clear signature of a state with octupolar correlations, although it does not strongly constrain the nature of these correlations to resemble those in spin ice.

Drawing conclusions about the nature of the magnetic state in $Ce_2Sn_2O_7$ is complicated significantly by the fact that measurements on samples grown by a different technique do not corroborate the observations of Ref. [219]. These new measurements do not observe significant scattering at large momentum transfer and instead point to dominantly dipolar correlations. Furthermore, as we describe in more detail the next section, these new results are consistent with constraints of the exchange parameters based on thermodynamic measurements, which suggest that $Ce_2Sn_2O_7$ has a long-range ordered all-in-all-out ground state, with an ordering temperature below the regime accessible by current experiments. The observed, spin-ice like dipolar correlations at low momentum transfer, hence, are a finite-temperature effect.

Since the two quite different conclusions in Ref. [219] and Ref. [223] are a consequence of different experimental data, clarification in this case with likely need more measurements, ideally on additional samples.

The results shown in Fig. 8.2 also suggest that single-crystal scattering would likely help to further constrain exchange parameters, since the angular dependence of scattering off the SI_x and the SI_y state is very different, enabling to distinguish the two even without a large angular resolution. This is important since neutron scattering measurements at large momentum transfer are challenging and one is unlikely to resolve finer features such as the pinch points.

8.2. Case for a $U(1)_\pi$ Quantum Spin Liquid Ground State in the Dipole-Octupole Pyrochlore $Ce_2Zr_2O_7$

The usual gold standard for fitting exchange parameters of magnetic materials is to compare inelastic neutron scattering data to predictions from spin wave theory. Sharp spin waves are expected to be the dominant excitation on top of most magnetically ordered states. Even when the ground state at zero field is not magnetically ordered, it is often possible to polarize the sample by means of an external field.

For dipolar-octupolar pyrochlores, the situation is more complicated because, as discussed in the previous section, neutron scattering at least at low momentum transfer is only sensitive to part of the pseudospin correlations. There are thus sets of parameters for which sharp spin waves are absent even in magnetically ordered

states, if the fluctuations are of purely octupolar character. The usual wisdom of equating the absence of sharp spin wave excitations with the absence of order, hence, does not extend to dipolar-octupolar pyrochlores. For this reason, it is necessary to develop alternative ways to constrain model parameters. In this section, we use simultaneous fits to the specific heat, magnetic susceptibility as well as neutron scattering to constrain the model parameters of $\text{Ce}_2\text{Zr}_2\text{O}_7$. Leaving uncertainties about the theoretical phase diagram aside, the results suggest that $\text{Ce}_2\text{Zr}_2\text{O}_7$ realizes an exotic $U(1)_\pi$ quantum spin liquid ground state.

The following people contributed to the experimental data shown in this section:

Sample Preparation: J. Dudemaine, E. M. Smith, J. Gaudet, A. Fitterman, J. Beare, C. A. Marjerrison, J. D. Garrett, G. M. Luke, B. D. Gaulin, and A. D. Bianchi. **Neutron scattering:** E.M. Smith, J. Gaudet, A. R. Wildes, S. Bhattacharya, C. R. C. Buhariwalla, N. P. Butch, J. P. Clancy, and E. Kermarrec. **Specific heat:** D. R. Yahne, T. DeLazzer, E. M. Smith, R. Movshovich, B. D. Gaulin and K. A. Ross. **Magnetic Susceptibility:** J. Beare, E. M. Smith, B. D. Gaulin, and G. M. Luke

8.2.1. Model Hamiltonian and fitting procedure

Analysis of the crystal field levels of $\text{Ce}_2\text{Zr}_2\text{O}_7$ show that its crystal electric field (CEF) ground state is given by a dipolar-octupolar doublet [220, 221], and to a good approximation this doublet is given by two $J = \frac{5}{2}$ states with $m_J = \pm\frac{3}{2}$. As discussed in detail in Sec. 7.3, the most general symmetry-allowed nearest-neighbor Hamiltonian for such a dipolar-octupolar doublet on the pyrochlore lattice is given by

$$H = H_{\text{ex}} + H_{\text{field}}, \quad (8.11a)$$

$$H_{\text{ex}} = \sum_{\langle ij \rangle} J_x S_i^x S_j^x + J_y S_i^y S_j^y + J_z S_i^z S_j^z, \quad (8.11b)$$

$$H_{\text{field}} = -\mu_B g_z \sum_j \mathbf{H} \cdot \hat{\mathbf{z}}_j (\sin \theta S_j^x + \cos \theta S_j^z). \quad (8.11c)$$

with four free parameters J_x , J_y , J_z , and θ . Furthermore, μ_B is the Bohr magneton and the g-factor g_z is fixed by the wave functions of the lowest CEF doublet, with $g_z = 2.57$ for a pure $J = \frac{5}{2}$ doublet [219–221]. The value in the real material might deviate slightly from this value due to mixing with states of higher total angular momentum J .

In the following, we outline the general fitting procedure used to determine the values of the four free parameters in Eq. (8.11) by fitting specific heat, magnetic susceptibility, and structure factor. To this end, it is useful to consider a permutation of the couplings (J_x, J_y, J_z) , as we already did when discussing the phase diagram of Eq. (8.11) in Sec. 7.3. Specifically, consider (J_a, J_b, J_c) such that $|J_a| > |J_b|, |J_c|$ and $J_b > J_c$, which allows to fit a subset of parameters using different measurements.

First, since the specific heat in the absence of an external magnetic field is insensitive to the permutation of the couplings, it should be fitted only within the subspace of parameter space spanned by the (J_a, J_b, J_c) . In particular, the largest coupling J_a simply plays the role of an overall energy scale and can be fitted from the high-temperature tail of the specific heat. The two other couplings are then most easily

parametrized by defining $J_{\pm} = -\frac{1}{4}(J_b + J_c)$ and $J_{\pm\pm} = \frac{1}{4}(J_b - J_c)$. In terms of these parameters, the constraints on J_a and J_b reduce the parameter space to a triangle as seen already in the phase diagram in Fig. 7.7. We cover this triangle uniformly by 2279 pairs $(J_{\pm}, J_{\pm\pm})$ and for each pair compute the specific heat as a function of temperature. Comparison to the experimental data then defines a goodness-of-fit

$$\left\langle \frac{\delta^2}{\epsilon^2} \right\rangle \propto \sum \frac{[C_{\text{mag}}^{\text{th}}(T_{\text{exp}}) - C_{\text{mag}}^{\text{exp}}(T_{\text{exp}})]^2}{\epsilon(T_{\text{exp}})^2}, \quad (8.12)$$

where $C_{\text{mag}}^{\text{th}}(T_{\text{exp}})$ and $C_{\text{mag}}^{\text{exp}}(T_{\text{exp}})$ are the theory and experimental result for the specific heat at temperature T_{exp} , respectively, and $\epsilon(T_{\text{exp}})$ is the uncertainty on the experimental data. Minimizing this quantity already constrains the possible values of the transverse exchange parameters $(J_{\pm}, J_{\pm\pm})$, although in practice the constraint is often not very tight (e.g. Fig. 8.3).

To fully specify the magnetic Hamiltonian in Eq. (8.11), it is necessary to also consider the form of the magnetic moment μ . This is determined by which permutation of (x, y, z) corresponds to (a, b, c) , as well as the value of the mixing angle $\theta = [0, \pi/2)$. This can be probed either by applying an external magnetic field (see chapter 9) or by measuring magnetic properties such as the susceptibility and the structure factor. In all cases, for measurements sensitive to the form of the magnetic moment experimental results have to be compared to theory computations using all three inequivalent permutations, $(x, y, z) = (a, b, c)$, $(x, y, z) = (b, c, a)$ and $(x, y, z) = (b, a, c)$ as well as a number of values for θ . In practice, the number of parameter sets that have to be considered can be reduced quite significantly by only considering those for which a good fit of the specific heat is possible.

In order to compare experimental measurements to theory calculations quantitatively, it is necessary to identify methods which are reliable for the model in its full parameter space, even if they are not controlled in any formal sense. Here, we use numerical linked cluster expansion (NLCE) as the main method for computing thermodynamic quantities at high temperature. The method in spirit is similar to high-temperature expansion but has often converges down to lower temperatures, see [252] for a pedagogical introduction. Additionally, it can be optimized quite significantly on the pyrochlore lattice [253]. All NLCE calculations shown in this thesis of specific heat and susceptibility have been performed by Robin Schäfer, and those of the structure factor by Owen Benton.

Since it is difficult to compute dynamical properties using NLCE, we also use semiclassical molecular dynamics simulations to compare to inelastic neutron scattering signatures. To this end, we first perform classical Monte Carlo simulations to obtain an ensemble of classical spin configurations $\{\mathbf{S}_j\}$ sampled at temperature T . These configurations are then used as initial configurations for numerical integration of the semiclassical Landau-Lifshitz equations

$$\frac{d}{dt}\mathbf{S}_j = -\mathbf{S}_j \times \mathbf{h}_j, \quad (8.13)$$

where \mathbf{h}_j is the effective magnetic field acting on site j . The classical dynamical structure factor $S_{\text{classical}}(\mathbf{q}, \omega)$ is obtained as the Fourier transform in space and time of the time-evolved magnetic moments, averaged over the ensemble of initial states. The semiclassical estimate for the dynamical structure factor of the quantum system

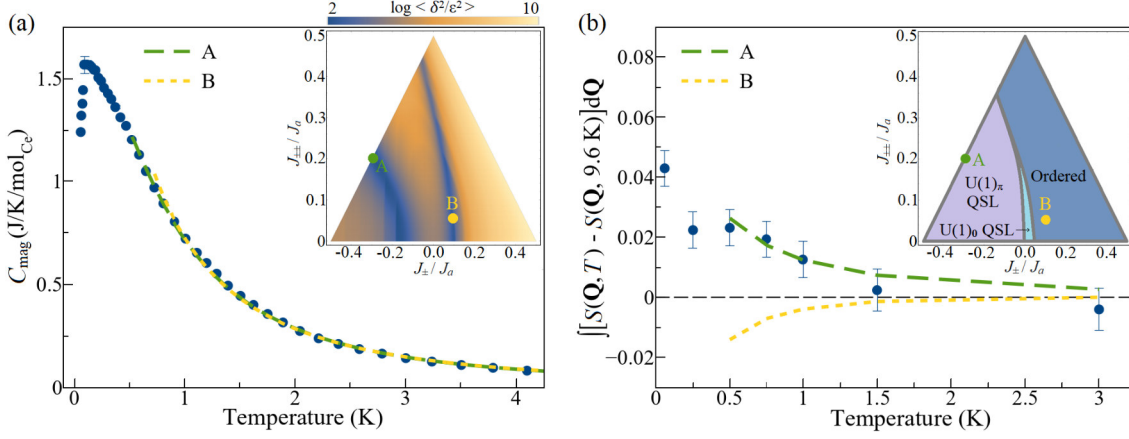


Figure 8.3.: Result of fitting specific heat (a) and integrated structure factor (b) for $\text{Ce}_2\text{Zr}_2\text{O}_7$. Blue markers in both panels show the experimental data, while dashed lines indicate theory calculations using NLCE. The inset in (a) shows the goodness-of-fit as a function of exchange parameters. Good fits lie in two extended regions, with the best parameter sets in each region indicated by A and B. Theory results for these two sets are indicated for both specific heat and structure factor. The figure is reprinted from Ref. [222].

is then obtained by rescaling its classical limit with an energy and temperature dependent factor

$$S_{\text{quantum}}(\mathbf{q}, \omega, T) = \frac{\beta\omega}{1 - e^{-\beta\omega}} S_{\text{classical}}(\mathbf{q}, \omega, T). \quad (8.14)$$

This factor can be motivated by comparing classical and quantum spin-wave theory at leading order [254], or by comparing the quantum and the classical version of the fluctuation-dissipation theorem [222]. For details on this, as well as numerical aspects of performing the simulations, see App. E.

8.2.2. Results

Placement in the phase diagram

The goodness-of-fit $\langle \delta^2/\epsilon^2 \rangle$ for the specific heat of $\text{Ce}_2\text{Zr}_2\text{O}_7$ is shown as a function of transverse exchange in the inset of Fig. 8.3 (a). It is minimized within two valleys, with the left valley lying completely within the $U(1)_\pi$ QSL phase and the right valley lying within the ordered phase. To determine the ground state of $\text{Ce}_2\text{Zr}_2\text{O}_7$, it is hence sufficient to determine which of these two regions of parameter space include its exchange parameters. In panel (a) of the figure, the two best fits within each region are marked as parameter sets A and B, and it also shows the specific heat as computed from fourth-order NLCE, together with the experimental data. In order to determine what parameter set better models $\text{Ce}_2\text{Zr}_2\text{O}_7$, panel (b) shows a comparison of the integrated structure factor $\int S(\mathbf{q})d\mathbf{q}$ computed from third-order NLCE and measured in experiment. Clearly, parameter set A compares well with the data while set B fails to capture the trend of the data even qualitatively. While the curves shown in the figure are for $\theta = 0$ and the permutation $(x, y, z) = (a, b, c)$, the trend of parameter set B that the structure factor decreases at low temperature

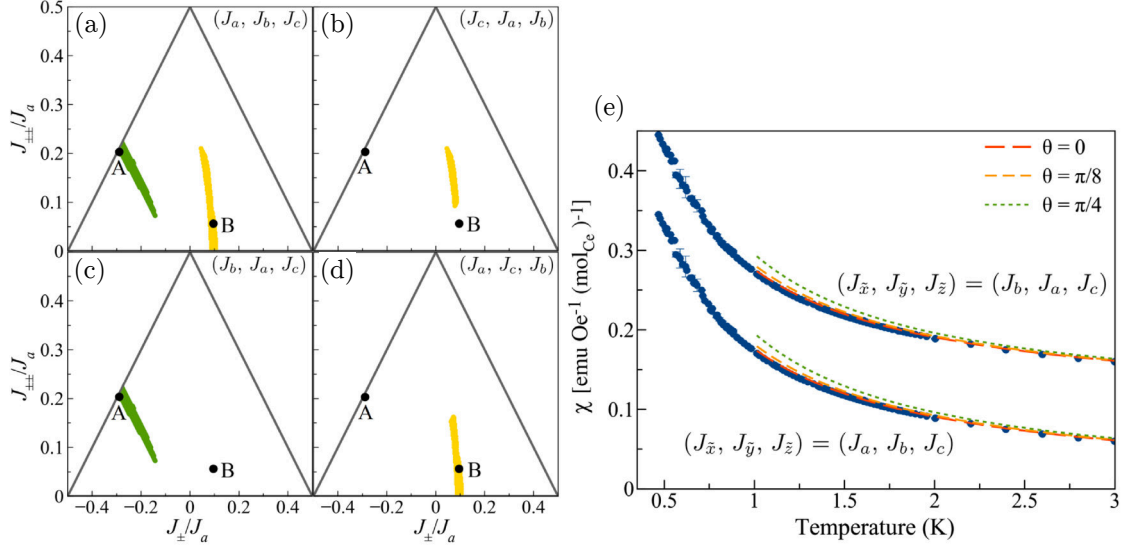


Figure 8.4.: Fits to susceptibility of $\text{Ce}_2\text{Zr}_2\text{O}_7$. Panel (a-d) show the regions of parameter space where a simultaneously good fit is achieved for the specific heat and susceptibility (see main text for details), for different permutation of the exchange parameters. A good fit for parameter region A is only possible for $J_c = J_z$. Panel (e) shows the susceptibility for different values of the mixing angle θ for these two permutations is shown. The data for one permutation is shifted upwards for better visibility. In both cases, the best fit is obtained for $\theta \approx 0$. The figure is reprinted from Ref. [222].

is true for all three permutations and all θ . The structure factor measurement hence does not tightly constrain the exact permutation of exchange parameters or the mixing angle, but it does rule out parameter set B conclusively.

In summary, we conclude that the measurements on $\text{Ce}_2\text{Zr}_2\text{O}_7$ are better described by set A, suggesting that this material realizes an exotic π -flux quantum spin ice state. The best-fit exchange parameters (set A) are

$$(J_a, J_b, J_c) = (0.064, 0.063, 0.011) \text{ meV}. \quad (8.15)$$

Nature of the local magnetic moment

To constrain the value of the mixing angle θ and fix the permutation of exchange parameters, we compare theory calculations using the best-fit parameters [Eq. (8.15)] to additional experimental measurements. The two main results are shown in Fig. 8.4 and Fig. 8.5, which contain fits of the magnetic susceptibility χ and the energy-resolved structure factor $S(\mathbf{q}, \omega)$, respectively.

Panels (a-d) of Fig. 8.4 show the regions of parameter space for which a good fit is achieved simultaneously for the specific heat, and the magnetic susceptibility. Both are modeled by NLCE and a “good fit” is defined as the goodness-of-fits being $\log\left(\left\langle\frac{\delta^2}{\epsilon^2}\right\rangle\right)_C < 2.7$ for the specific heat and $\log\left(\left\langle\frac{\delta^2}{\epsilon^2}\right\rangle\right)_\chi < 12.1$ for the magnetic susceptibility. As is shown in the figure, only the two permutations with $J_c = J_z$ allow for a simultaneously good fit within the region of parameter space containing




Figure 8.5.: Elastic neutron scattering from a powder sample of $\text{Ce}_2\text{Zr}_2\text{O}_7$ as a function of temperature. Panel (a-c) show experimental data and panel (d-f) and (g-i) show the result of semiclassical molecular dynamics calculations for $\theta = 0$ and $\theta = \pi/2$ respectively. The figure is reprinted from Ref. [222].

set A . Since set B has been ruled out based on the integrated structure factor, panel (e) of the figure shows the calculated magnetic susceptibility for these two permutations for different values of the mixing angle. In both cases, the best fit is achieved by a mixing angle close to zero $\theta \approx 0$. Note that since $J_a \approx J_b$, the two permutations $(x, y, z) = (a, b, c)$ and $(x, y, z) = (b, a, c)$ lead to almost identical exchange Hamiltonians.

The result $\theta \approx 0$ is consistent also with the result of inelastic neutron scattering measurements performed on powder samples of $\text{Ce}_2\text{Zr}_2\text{O}_7$. These results are shown in Fig. 8.5, together with theory calculations for $\theta = 0$ and $\theta = \pi/2$ obtained from semiclassical molecular dynamics simulations. Increasing θ from zero shifts spectral weight to lower energies, which is inconsistent with the measured distribution. The simulation also captures the temperature dependence of the signal at least qualitatively.

Overall best parameters

In summary, we conclude that all experimental measurements are consistent with parameters

$$(J_x, J_y, J_z) = (0.064, 0.063, 0.011) \text{ meV} \quad \text{and} \quad \theta \approx 0. \quad (8.16)$$

Importantly, the above parameter set should be understood as the representative of what is really a family of good-fit parameters. This is because the specific heat fit presented in Fig. 8.3 (a) does not tightly constrain the exact value of the exchange parameters, and this remains true when also fitting the susceptibility, as shown in panels (a,c) of Fig. 8.4. Equally important is the observation that even within this

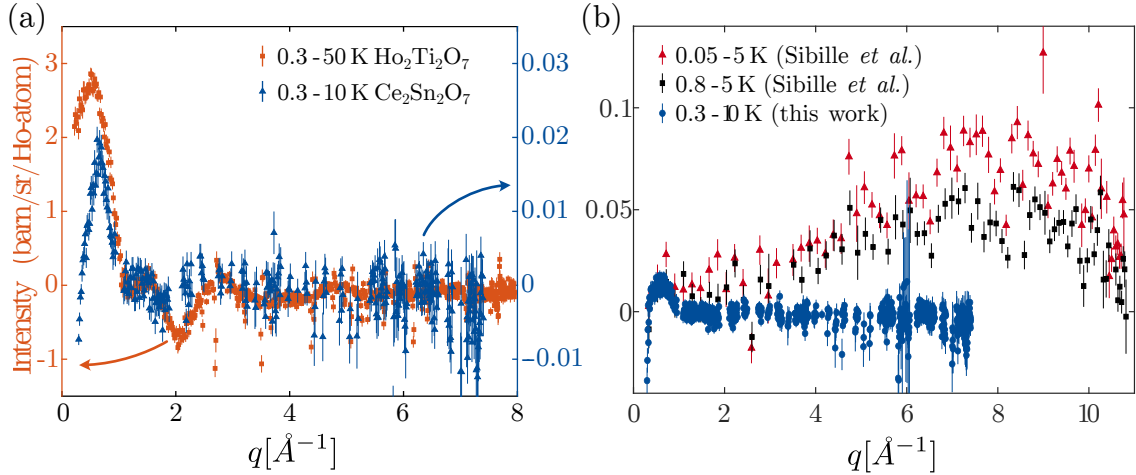


Figure 8.6.: Neutron scattering cross section of $\text{Ce}_2\text{Sn}_2\text{O}_7$ samples grown using hydrothermal techniques. Panel (a) compares the cross section to that of the classical spin ice $\text{Ho}_2\text{Ti}_2\text{O}_7$, panel (b) to the cross section as measured on solid-state samples (Sibille *et al.* [219]).

residual uncertainty, the material is expected to realize an exotic $U(1)_\pi$ quantum spin ice ground state, with mostly octupolar correlations.

The fitted exchange Hamiltonian also for an intuitive explanation for the observed absence of sharp spin wave excitations even in field-polarized samples of $\text{Ce}_2\text{Zr}_2\text{O}_7$. Since the field-polarized state is polarized in the pseudospin- z component and $\theta \approx 0$, transverse fluctuations carry only an octupolar moment and do not couple to neutrons at least at low momentum transfer.

Note that the results presented here are fully consistent with the results of an independent study presented in Ref. [226].

8.3. Dipolar Spin Ice Regime Proximate to an All-In-All-Out Néel Ground State in the Dipolar-Octupolar Pyrochlore $\text{Ce}_2\text{Sn}_2\text{O}_7$

The material $\text{Ce}_2\text{Sn}_2\text{O}_7$ is a close sibling of $\text{Ce}_2\text{Zr}_2\text{O}_7$ with the transition metal zirconium substituted by tin. As mentioned already briefly in Sec. 8.1, it has received much attention because of the study in Ref. [219], which reports observation of a neutron scattering cross section peaked at large momentum transfer in a powder sample. Such scattering is a direct signature of octupolar correlations, and in the absence of sharp features in both specific heat and susceptibility the authors conclude that the sample realizes a quantum liquid of magnetic octupoles. Based on fitting a simplified model Hamiltonian and using mean field calculations, the sample is also placed in the phase diagram within the $U(1)_\pi$ QSL phase close to the classical spin ice point. This is marked by a black square ($\text{Ce}_2\text{Sn}_2\text{O}_7$, Sibille *et al.*) in Fig. 7.7.

In this chapter, we fit the exchange Hamiltonian $\text{Ce}_2\text{Sn}_2\text{O}_7$ using a more sophisticated setup, similar to the one used for $\text{Ce}_2\text{Zr}_2\text{O}_7$ in the last chapter. This is motivated by new experimental measurement results obtained from samples grown by hydrothermal crystal growth [255]. This synthesis technique is different from the

solid state techniques used for sample growth in previous studies [218, 219]. The main difference of hydrothermal synthesis is the significantly lower temperature required (700°C), as compared with solid-state synthesis (typically $> 1000^\circ\text{C}$). This avoids temperatures at which tin oxide becomes volatile, thereby ideally minimizing oxidation. Oxidation is an important issue in Cerium pyrochlores since both its trivalent Ce^{3+} and tetravalent Ce^{4+} oxidation state are stable, the latter being nonmagnetic.

Strikingly, neutron scattering measurements on these new samples do not confirm the existence of scattering at large momentum transfer. Instead, the cross section resembles that of the dipolar spin ice $\text{Ho}_2\text{Ti}_2\text{O}_7$, as shown in Fig. 8.6. Regardless of what is the reason for the different results on these two samples, it is desirable to see whether the new results allow for a consistent description when including also specific heat and susceptibility measurements. Indeed, we find a remarkably good agreement between theory and experiment for all three measured quantities using the parameters

$$(J_x, J_y, J_z) = (0.045, -0.001, -0.012) \text{ meV}, \quad \theta = 0.19\pi, \quad g_z = 2.2. \quad (8.17)$$

This places the material in the ordered phase (see Fig. 7.7, Yahne et al.) close to the phase boundary to the $U(1)_0$ quantum spin liquid. Hence, our results suggest that $\text{Ce}_2\text{Sn}_2\text{O}_7$ has a magnetically ordered dipolar all-in-all-out ground state, with an ordering temperature below the regime currently accessible by experiments.

As will be explained in the remainder of this section, we also improve on the general fitting procedure as described in the previous section for $\text{Ce}_2\text{Zr}_2\text{O}_7$ in multiple ways. First, we here also fit the value of the g-factor g_z , to account for possible mixing of the crystal field ground state with states of total angular momentum $J = \frac{7}{2}$. Second, we perform a full parameter sweep for three quantities, the specific heat, magnetic susceptibility, and the powder-averaged structure factor. For the latter sweep, we use classical Monte-Carlo simulations, which we expect to give a reasonably good description of the experimental results given the relatively large temperature $Tk_B \approx 0.5J_x$. Note that, since the measurement is also conducted over a large range of momentum transfer it is important to use the full dipolar-octupolar form factor (see Sec. 8.1) beyond the dipole approximation. Finally, the parameters in Eq. (8.17) allow for efficient numerical solution of the dipolar-octupolar Hamiltonian using Quantum Monte-Carlo. Because of this, we also implement the stochastic series expansion to refine the specific heat fit around the best-fit value found from the full parameter sweep computed using NLCE. This also allows us to estimate the critical temperature of the transition into the ordered ground state regime.

The following people contributed to the experimental data shown in this section: **Sample preparation:** M. Powell and J. W. Kolis. **Neutron scattering:** D. R. Yahne, S. Calder, and M. D. Frontzek. **Specific heat:** D. R. Yahne, A. F. May, and C. M. Pasco. **Magnetic Susceptibility:** A.F. May and C. M. Pasco.

8.3.1. Specific heat

One marked success of the hydrothermal synthesis technique is the growth of single crystals, one example shown in the inset of Fig. 8.7 (a). Importantly, this makes it possible to make specific heat measurements down to low temperature

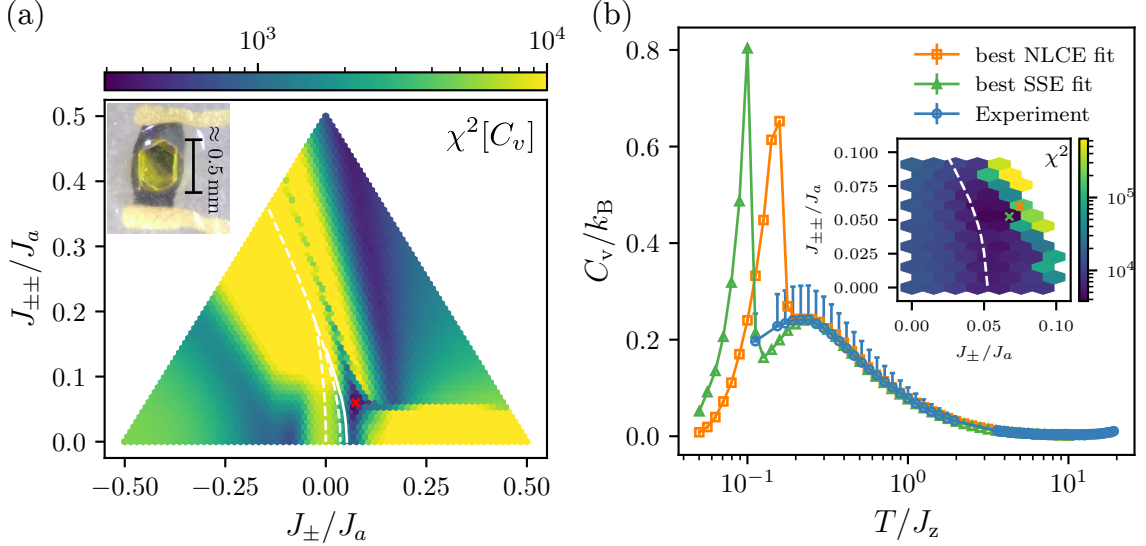


Figure 8.7.: Parameter fit to the specific heat C_v of $\text{Ce}_2\text{Sn}_2\text{O}_7$. Panel (a) shows the goodness-of-fit as a function of the transverse exchange couplings as obtained from the full parameter sweep using NLCE. Panel (b) shows the results of a refined parameter fit, from a sweep around the best-fit value of the NLCE sweep using quantum Monte-Carlo.

(~ 50 mK, while previous measurements on powder samples were limited to temperatures $\gtrsim 0.4$ K [219]). We use these extended measurements to constrain the exchange parameters of $\text{Ce}_2\text{Sn}_2\text{O}_7$ using an almost identical procedure as used for $\text{Ce}_2\text{Zr}_2\text{O}_7$ in the previous section.

We extract for each pair of transverse exchange parameters ($J_\pm, J_{\pm\pm}$) the overall energy scale J_a by fitting the high-temperature tail of the specific heat. Because the sample is very small, its mass comes with a significant uncertainty which enters the specific heat measurement as a temperature dependent systematic error $\delta_{\text{sys}}(T)$. Because of this, we fit the mass of the sample in addition to the exchange parameters, so that for each ($J_\pm, J_{\pm\pm}$) we minimize the goodness-of-fit

$$\chi^2[C] = \sum_{0.1\text{ K} < T_i < 1.9\text{ K}} \frac{[C_{\text{NLCE}}(T_i) - C_{\text{exp}}(T_i) - \varepsilon \delta_{\text{sys}}(T_i)]^2}{\delta C_{\text{NLCE}}(T_i)^2 + \delta_{\text{stat}}(T_i)^2} \quad (8.18)$$

with respect to $\varepsilon \in [-1, 1]$. Above, C_{NLCE} and C_{exp} denote the specific heat as computed using NLCE and as measured in experiment, respectively, and δC_{NLCE} and δ_{stat} are the respective error estimates. The NLCE error is estimated using the difference between successive orders (see App. B of Ref. [223] for details).

The goodness-of-fit as a function of the transverse exchange parameters is shown in Fig. 8.7 (a). Remarkably, it is much more tightly constrained than the same fit for $\text{Ce}_2\text{Zr}_2\text{O}_7$ (cf. Fig. 8.3), despite the additional fitting parameter ε . Moreover, the best fit parameters from the NLCE fit

$$(J_a, J_b, J_c)_{\text{NLCE}} = (0.0452, -0.0014, -0.0122) \text{ meV} \quad (8.19)$$

are quite different from the estimate for $\text{Ce}_2\text{Zr}_2\text{O}_7$. This is somewhat surprising given their chemical similarity, but the result hinges on the fact that the measured specific heat of the two compounds is quite different.




Figure 8.8.: Parameter fit to the magnetic susceptibility χ of $\text{Ce}_2\text{Sn}_2\text{O}_7$. Panel (a) shows the goodness-of-fit for the best specific heat fitting parameters [Eq. (8.19)], as a function of the mixing angle θ for the three inequivalent permutations of exchange parameters. Panel (b) shows the experimental data together with the best fit for each permutation.

The phase boundaries from cluster mean field theory and quantum Monte-Carlo are indicated in Fig. 8.7 (a) by white dashed and solid lines, respectively. Clearly, all good fit parameters lie in the ordered all-in-all-out phase and they admit sign-problem-free simulation using quantum Monte-Carlo, since both transverse exchanges J_b and J_c are ferromagnetic. Since no singularity was observed in the measured specific heat, the transition temperature to the ordered state must be lower than the regime accessible by experiment.

To get an estimate of the critical temperature, we implement the stochastic series expansion [256] with the extended cell update [212] (see Ref. [257] for a pedagogical review). As is shown in Fig. 8.7 (b), the critical temperature for the best NLCE fit parameters is above the lowest measured temperatures and they are thus incompatible with the experiment. Since the parameters are very close to the phase boundary to the $U(1)_0$ state, the transition temperature is extraordinarily sensitive to the exact value of the exchange parameters. Because of this, we refit the experimental results using QMC in a region around the best NLCE fit. The resulting goodness-of-fit is shown in the inset of Fig. 8.7 (b). A very slight change of exchange parameters by about 1% to

$$(J_a, J_b, J_c)_{\text{NLCE}} = (0.0457, -0.0014, -0.0110) \text{ meV} \quad (8.20)$$

suffices to push the transition below the lowest temperatures measured in experiment, without sacrificing any accuracy at high temperatures. For all good fits, the fitting parameter for the sample mass is close to one $\varepsilon \approx -1$, meaning that obtaining a good fit requires a samples mass at the limit of that consistent with the measurements. This is also indicated in Fig. 8.7 (b), where we shift the experimental data points (blue) within the error bars marking the systematic uncertainty.

8.3.2. Magnetic susceptibility

While the fitting of the specific heat discussed above puts reasonably tight constraints on the placement of $\text{Ce}_2\text{Sn}_2\text{O}_7$ in the phase digram of dipolar-octupolar

pyrochlores, it is insensitive to the form of the magnetic moment. To constrain the form of the moment, we also fit measurements of the magnetic susceptibility. In Fig. 8.8 (a), we show the goodness-of-fit for the best specific heat fit parameters [Eq. (8.19)] as a function of the mixing angle θ and for the three inequivalent permutations of exchange parameters. For each parameter set we also fit the value of the g-factor g_z . The calculated specific heat in all cases is obtained using NLCE.

Although we performed a full parameter sweep, only the best specific heat fit parameters are shown since good fits to the susceptibility can be obtained for a wide range of exchange parameters. As is clear from Fig. 8.8, the susceptibility is, however, quite sensitive to the mixing angle θ , with sharp minima for two of the three inequivalent permutations. While the global goodness-of-fit minimum is achieved for the permutation $(J_x, J_y, J_z) = (J_b, J_c, J_a)$, no good fit of the structure factor is possible in this case, while taking $(J_x, J_y, J_z) = (J_a, J_b, J_c)$ admits a good fit (see below).

8.3.3. Structure factor

The neutron scattering measurements performed on the samples grown using hydrothermal techniques do not agree with previous results obtained from samples grown via solid state synthesis [219]. It is therefore important to verify that the new results can be described consistently with the thermodynamic measurements discussed above. To this end, we perform another full parameter sweep, using the same parametrization as for the susceptibility, but computing for each parameter set the powder-averaged structure factor.

More specifically, we employ classical Monte-Carlo simulations to compute the powder-averaged structure factor $S(|\mathbf{q}|)$. In contrast to the specific heat, the structure factor depends on the permutation of exchange parameters, that is on a specific choice of (J_x, J_y, J_z) as well as on the value of the mixing angle θ . Because of this, we can use it to fit the full set of these parameters. To this end, for each of the 2279 pairs of $(J_{\pm}, J_{\pm\pm})$, we first run a classical Monte-Carlo simulation of the XYZ Hamiltonian Eq. (7.27) at $T = 10$ K and $T = 0.3$ K. We use the energy scale α obtained from a NLCE fit to the high-temperature tail of the specific heat to obtain scaled exchange parameters

$$J_a = \alpha, \quad J_b = 2\alpha(J_{\pm} + J_{\pm\pm}), \quad J_c = 2\alpha(J_{\pm} - J_{\pm\pm}). \quad (8.21)$$

In each simulation, we compute the correlation functions

$$C_{ij}^{\alpha\beta}(\mathbf{q}) = \frac{1}{N} \sum_{\mathbf{r}, \mathbf{r}'} \exp[-i\mathbf{q} \cdot (\mathbf{r} - \mathbf{r}')] \langle S_i^{\alpha}(\mathbf{r}) S_j^{\beta}(\mathbf{r}') \rangle \quad (8.22)$$

where $S_i^{\alpha}(\mathbf{r})$ is the $\alpha \in \{a, b, c\}$ component of a classical Heisenberg spin on sublattice $i \in \{0, 1, 2, 3\}$ in the unit cell at position \mathbf{r} . Different permutations of the exchange parameters as well as different choices of θ are then implemented by a suitable rotation and permutation of components of the correlation function $C_{ij}^{\alpha\beta}(\mathbf{q}) \rightarrow C_{ij}^{\tilde{\alpha}\tilde{\beta}}(\mathbf{q})$, without the need to run a new Monte-Carlo simulation. We then follow Sec. 8.1 to compute the full dipolar-octupolar structure factor $S(\mathbf{q})$ for all three inequivalent permutations of exchange parameters as well as for 101 values of θ in $[0, \pi/2]$.

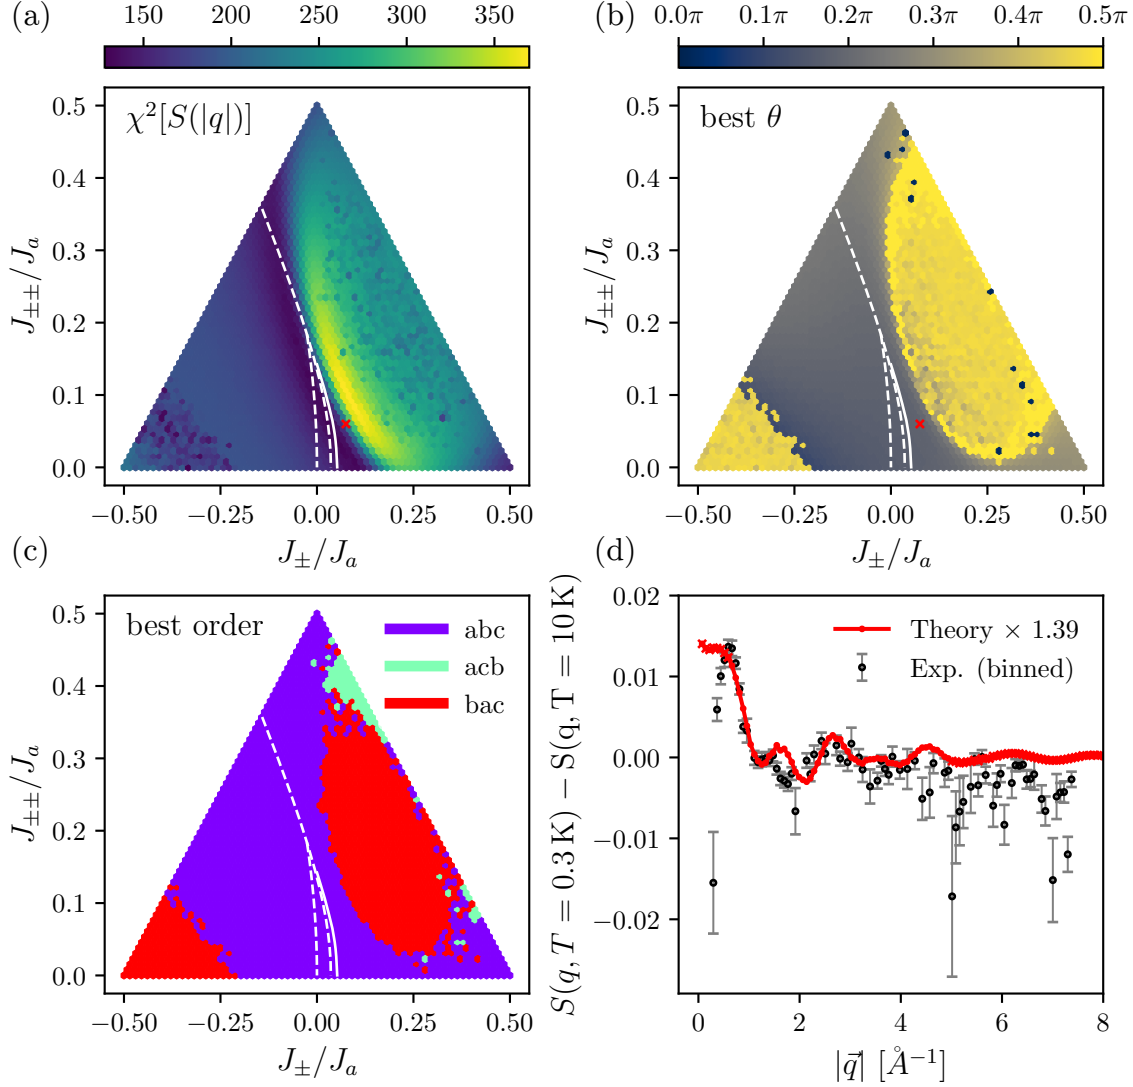


Figure 8.9.: Parameter fit to the powder-averaged structure $S(|\mathbf{q}|)$ of $\text{Ce}_2\text{Sn}_2\text{O}_7$. Panel (a) shows the goodness-of-fit χ^2 as a function of the transverse exchange couplings as obtained from the full parameter sweep using classical Monte-Carlo. For each pair of $(J_{\pm}, J_{\pm\pm})$, χ^2 is minimized with respect to the mixing angle θ and the three inequivalent permutations of exchange parameters. Panels (b) and (c) show the best θ and best permutation as a function of the transverse exchange, respectively. Panel (d) shows the (binned) experimental data together with the parameters for the best combined fit [Eq. (8.17)].

Finally, we average over direction to obtain the powder-averaged structure factor $S(|\mathbf{q}|)$. Since our simulation are performed on a cluster of finite size, we only sample a discrete set of momenta \mathbf{q} . To perform the powder average, we bin these into momentum shells of width $dq = 2\pi/(8a)$, where a is the lattice constant, and average over these shells. To compute the goodness-of-fit, we also rebin the experimental data into the same shells and define

$$\chi^2 [S(|\mathbf{q}|)] = \sum_{q_{\min} < |\mathbf{q}| < q_{\max}} \frac{(\lambda S_{\text{MC}}(|\mathbf{q}|) - S_{\text{exp rebinned}}(|\mathbf{q}|))^2}{\delta_{\text{exp rebinned}}^2(|\mathbf{q}|)}, \quad (8.23)$$

where λ is an additional scale parameter that is minimized for each dataset as a consistency check and $q_{\min} = 0.66 \text{ \AA}^{-1}$ and $q_{\max} = 5 \text{ \AA}^{-1}$. Since we are comparing absolute differential cross sections, λ for good fits should be close to one. Indeed we find that for the best parameter estimate [Eq. (8.17)], $\lambda = 1.39$. Note that the exact value of this scale factor λ is quite sensitive to the mixing angle θ . To reproduce the exact value of $\lambda = 1.39$ as a best fit, one has to choose the exact best fit value of $\theta = 0.185\pi \approx 0.19\pi$. In Fig. 8.9 (a), we show the goodness-of-fit, optimized over the three inequivalent permutations of exchange parameters and all values of the mixing angle θ . Panel (b) and (c) of the same figure show the optimal values of the mixing angle θ and the optimal permutation for each pair of transverse exchange, respectively. Remarkably, for all good fits, the optimal permutation of exchange parameters is $(J_x, J_y, J_z) = (J_a, J_b, J_c)$, and the optimal value of the mixing angle is $\theta \approx 0.2\pi$. This independent fit is thus consistent with the result for the magnetic susceptibility, which also yields $\theta \approx 0.2\pi$. Also, the optimal scale λ is of order one for all good fits. In summary, while the fit to the structure factor does constrain the exchange permutation and the value of the mixing angle quite reliably, it does not strongly constrain the value of exchange parameters.

8.3.4. Overall best fit

The overall best fit is obtained by minimizing the combined goodness of fit of specific heat and structure factor

$$\chi_{\text{tot}}^2 = \chi^2[S(|\mathbf{q}|)] + \chi^2[C_v(T)]. \quad (8.24)$$

This yields the values quoted earlier [Eq. (8.17)] and is marked in Fig. 8.7 (a) and Fig. 8.9 (a) by a red cross. Note that since the fit to the specific heat puts much tighter constraints on the exchange parameters compared to that to the structure factor, the combined best fit is very close to the best fit to the specific heat. The structure factor for the best combined fit it shown in Fig. 8.9 (d), together with the experimental data. Considering also the best fit of the specific heat in Fig. 8.7 and the best fit of the magnetic susceptibility in Fig. 8.8, we obtain a good description for three independent measurement within a single model with only six free parameters.

8.4. Conclusion and Outlook

In summary, in this chapter we shown that it possible to obtain model parameters of exchange Hamiltonians for dipolar-octupolar pyrochlores even in the absence of visible spin wave excitations, using simultaneous fits to multiple thermodynamic

measurements. After explaining the complications leading to the absence of spin waves in dipolar octupolar pyrochlores in Sec. 8.1, we described the general idea and procedure as originally developed for $\text{Ce}_2\text{Zr}_2\text{O}_7$. In this case we used fits to the heat capacity and the integrated structure factor to constrain the exchange couplings of the material to lie within a region of parameter space that is believed to host an exotic $U(1)_\pi$ quantum spin liquid ground state. Subsequent fits to the magnetic susceptibility and the dynamical structure factor revealed that correlations are most likely octupolar in nature, implying the absence of sharp, observable spin wave excitations even in field polarized samples. This demonstrates both, the necessity and the success of fitting to multiple measurements for reliable parameter estimation.

In Sec. 8.3, we showed the results of applying a similar, but significantly refined, fitting program to another Cerium-based pyrochlore, $\text{Ce}_2\text{Sn}_2\text{O}_7$. Remarkably, the fits suggest that $\text{Ce}_2\text{Sn}_2\text{O}_7$ has a ground state with dipolar all-in-all-out order, but its parameters in the phase diagram are located close to the phase boundary with the $U(1)_0$ liquid phase. The lack of sharp features in specific heat, susceptibility, and structure factor are hence a consequence of the very low critical temperature. The critical temperature lies below the regime accessed by current experiments, as we showed explicitly using quantum Monte-Carlo simulations with the best-fitting parameters. Because the proximity to the phase boundary, the transition temperature is extremely sensitive to the exact value of exchange parameters, making it impossible to provide reliable prediction for future experiments. It does show, however, that the absence of a transition in the experiment is consistent with exchange parameters very close to the best fit. Fits to the susceptibility and the structure factor are also consistent with these parameters and yield consistent and independent estimates of the form of the magnetic moment. Correlations are shown to be dominantly dipolar in nature, consistent with the fact that the neutron scattering cross section mimics that of the dipolar spin ice $\text{Ho}_2\text{Ti}_2\text{O}_7$.

For both materials, future work is needed to conclusively determine the exact nature of the ground state. In the case of $\text{Ce}_2\text{Zr}_2\text{O}_7$, an interesting possibility would be measuring neutron scattering cross section at large momentum transfer, similar to what was done on powder samples of $\text{Ce}_2\text{Sn}_2\text{O}_7$. Since large single crystals are available in this case, such measurement could even be attempted with some angular resolution. This would provide valuable insights into the nature of the correlations and in particular the anisotropy of the exchange Hamiltonian as explained in Sec. 8.1. For $\text{Ce}_2\text{Sn}_2\text{O}_7$, clarification is needed regarding the seemingly incompatible results reported here and in Ref. [219]. Possible explanations include an issue with one of the experiments, but also strong sample dependence caused, for example, by disorder. The effect of disorder in general is an important issue, especially in cerium based pyrochlores because of the existence of the stable but non-magnetic Ce^{4+} oxidation state. Such oxidized ions act effectively as magnetic vacancies and the effect of such on pyrochlores especially close to phase boundaries is not well understood but could well prove essential for their successful, quantitative modeling.

9. Field-Induced Spin Chains in Dipolar Octupolar Pyrochlores

Parts of the work presented in this chapter are also part of the following articles:

[248] Benedikt Placke, Roderich Moessner, and Owen Benton. “Hierarchy of energy scales and field-tunable order by disorder in dipolar-octupolar pyrochlores”. In: *Phys. Rev. B* 102 (24 Dec. 2020), p. 245102. DOI: 10.1103/PhysRevB.102.245102

[258] E. M. Smith, J. Dudemaine, B. Placke, et al. *Quantum Spin Ice Response to a Magnetic Field in the Dipole-Octupole Pyrochlore $Ce_2Zr_2O_7$* . arXiv:2306.13183. DOI: 10.48550/arXiv.2306.13183

In this chapter, we consider the effect that applying an external field has on dipolar-octupolar pyrochlore materials. We focus on the case of a strong external magnetic field applied in the [110] direction, in which case dipolar-octupolar pyrochlores are known to form a ‘chain’ state, with subextensive degeneracy. Magnetic moments are correlated along one-dimensional chains carrying effective Ising degrees of freedom which are noninteracting on the mean-field level. Here, we investigate this phenomenon in detail, including the effects of quantum fluctuations. We identify two distinct types of chain phases, both featuring distinct, subextensive, classical ground state degeneracy. Focussing on one of the two kinds, we discuss lifting of the classical degeneracy by quantum fluctuations. We map out the ground-state phase diagram as a function of the exchange couplings, using linear spin wave theory and real-space perturbation theory. We find a hierarchy of energy scales in the ground state selection, with the effective dimensionality of the system varying in an intricate way as the hierarchy is descended. We derive an effective two-dimensional anisotropic triangular lattice Ising model with only three free parameters which accounts for the observed behavior.

Connecting our results to experiment, they are consistent with the observation of a disordered chain state in $Nd_2Zr_2O_7$. We also show that the presence of two distinct types of chain phases has consequences for the field-induced breakdown of $U(1)$ octupolar quantum liquid phases possibly realized in $Ce_2Sn_2O_7$ and $Ce_2Zr_2O_7$. Finally, we compare theory calculations with new neutron scattering data of $Ce_2Zr_2O_7$. Confirming our earlier proposal (see Sec. 8.2), we find that correlations are mostly octupolar for all parameters compatible with the experimental results. We also identify direct signatures that would enable a more refined fitting of the anisotropy of the exchange interaction if the experimental resolution was to improve slightly.

The rest of the chapter is organized as follows. In Sec. 9.1, we review the classical ground-state phase diagram of dipolar-octupolar pyrochlores in a [110] field. This serves as a starting point for the order-by-disorder calculation presented in Sec. 9.2. In Sec. 9.2.1, we derive a two-dimensional triangular lattice Ising model describing

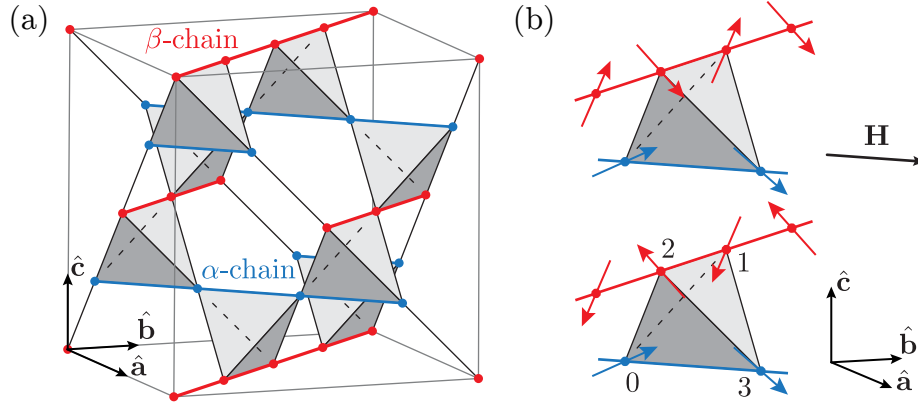


Figure 9.1.: (a) Pyrochlore lattice separated into chains parallel (α chains) and perpendicular (β chains) to the external field $\mathbf{H} \parallel (1, 1, 0)$. (b) The two classical ground-state orientations for a single tetrahedron with the local easy axis along the local $[111]$ direction. While the magnetic moments on the α chains are pinned by the external field, those on the β chains are perpendicular to the field and thus each retain one independent Ising-like degree of freedom. On one of the tetrahedra, we also indicate the four fcc sublattices 0, 1, 2, 3. The local magnetic moments on each sublattice point into a direction $\pm \hat{z}_i$, where the \hat{z}_i are defined to point into the tetrahedron.

the lifting of the classical degeneracy in leading-order real space perturbation theory (RSPT). In Sec. 9.2.2, we compute the ground-state phase diagram using linear spin wave theory (LSWT), comparing the results with those from RSPT. In Sec. 9.2.3 we discuss the tunability of the ground state selection as a function of the external field strength. Finally, sections 9.3 and 9.4 explore the connection of our work to experiment. We conclude in Sec. 9.5.

9.1. Mean-field Phase Diagram in $[110]$ Field

As a basis to study the effect of quantum fluctuation, we discuss in this section the classical mean field phase diagram of the exchange Hamiltonian of dipolar-octupolar pyrochlores as derived in Sec. 7.3

$$H = H_{\text{ex}} + H_{\text{field}}, \quad (9.1a)$$

$$H_{\text{ex}} = \sum_{\langle ij \rangle} J_x S_i^x S_j^x + J_y S_i^y S_j^y + J_z S_i^z S_j^z, \quad (9.1b)$$

$$H_{\text{field}} = -\mu_B g_z \sum_j \mathbf{H} \cdot \hat{z}_j (\sin \theta S_j^x + \cos \theta S_j^z). \quad (9.1c)$$

9.1.1. Zero field limit

Before turning to the case of a finite field, we briefly discuss the zero-field ($\mathbf{H} = 0$) phase diagram. In contrast to the quantum case discussed in Sec. 7.3.4, in the absence of transverse fluctuations there is no difference between the zero and π

flux state and instead both correspond to the same classical spin ice state. When distinguishing all three exchange parameters, there are then six different phases in total. These are the three ice-like phases SI_λ ($\lambda = x, y, z$), stabilized when the corresponding exchange coefficient is positive and sufficiently strong

$$J_\lambda > \max(-3J_{\lambda'}, J_{\lambda'}) \quad \forall \lambda' \neq \lambda. \quad (9.2)$$

In these phases, all pseudospins align with the λ axis and are governed by a local “two up, two down” constraint with respect to this axis on every tetrahedron. This rule leaves an extensive number of degrees of freedom. Assuming in the following generic $\theta \neq 0$, if $\lambda = x, z$ then the rule corresponds to a “two-in-two-out” rule on the magnetic moments; i.e. for each tetrahedron, two of the four magnetic moments on its corners point into and two point out of it. SI_x and SI_z thus describe the spin ice state already discussed in Sec. 7.1. In contrast, SI_y describes a state analogous to spin ice but where the active degrees of freedom are octupolar moments (see Sec. 7.3 for details).

There are then three ordered, “all-in-all-out” phases (AIAO_λ , $\lambda = x, y, z$), occurring for sufficiently strong negative exchange coefficient

$$J_\lambda < \min(-J_{\lambda'}/3, J_{\lambda'}) \quad \forall \lambda' \neq \lambda. \quad (9.3)$$

These correspond to ferromagnetic order of the pseudospins along the λ axis. If $\lambda = x, z$ then this will correspond to antiferromagnetic all-in-all-out order of the magnetic moments. The AIAO_y phase, by contrast, is a form of octupolar order.

9.1.2. Polarization of dipolar spin ice and ordered phases

Turning to finite fields, for the rest of this chapter we set

$$\mathbf{H} = \frac{\sqrt{3}h}{2\mu_B g_z} \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}. \quad (9.4)$$

Substituting this into Eq. (9.1) yields:

$$H_{\text{field}} = -h \sum_{i \in \mathcal{L}_0} (\sin \theta S_i^x + \cos \theta S_i^z) + h \sum_{j \in \mathcal{L}_3} (\sin \theta S_j^x + \cos \theta S_j^z), \quad (9.5)$$

where \mathcal{L}_ν ($\nu = 0, 1, 2, 3$) denotes the four fcc sublattices, also indicated in Fig. 9.1 (b). The field, when applied in this direction, only couples to sites on the sublattices 0 and 3. As shown in Fig. 9.1 (a), sites on which the magnetic moment couples to the external field then form one-dimensional chains which are parallel to the field direction (α chains), whereas sites whose magnetic moment does not couple to the external field also form one-dimensional chains, which are perpendicular to the field direction (β chains).

The classical ground state phase diagram of dipolar-octupolar pyrochlores with $\mathbf{H} \parallel (1, 1, 0)$ is shown in Fig. 9.2. There are now seven possible phases, six of which are directly connected to the six phases possible at zero field. The SI_Y phase evolves into two distinct chain phases as a function of field strength, separated by a second-order transition.




Figure 9.2.: Zero-temperature mean-field phase diagram of dipolar-octupolar pyrochlores [Eq. (9.1)] with an applied field along the [110] direction as a function of exchange couplings and field strength, for $J_z > 0$. The first column is for fixed octupolar coupling $J_y = 0$ and the second and third column are for fixed magnetic fields $h = J_z/2$ and $h = 4J_z$, respectively. The top and bottom rows have a fixed mixing angle $\theta = \pi/4$ and $\theta = 0.33\pi$. The phase diagram does not change qualitatively for any value of θ except for the singular limits $\theta \rightarrow 0$, $\theta \rightarrow \pi/2$, and $\theta \rightarrow \pi$. As $h \rightarrow 0$ (not shown), while the low-field all-in-all-out (AIAO $_\lambda$) phases are stable, the chain phases (Chain $_\lambda^{(*)}$) are not. In the dipolar case ($\lambda = x, z$), at exactly zero field the chain phases become the respective spin-ice phase (SI $_\lambda$), with extensive instead of subextensive degeneracy. For the octupolar chain phase, as the field is lowered starting in the Chain $_\gamma$ phase, there is actually a second-order transition to an intermediate phase, Chain $_\gamma^*$ [see also Fig. 9.3], in which each α chain carries an additional Ising degree of freedom. The Chain $_\gamma^*$ phase is connected to the octupolar ice phase SI $_\gamma$ at zero field. At large fields, the orientation of magnetic moments in the AIAO $_\lambda$ phases get significantly distorted with respect to their zero field orientations. However, even as $h \rightarrow \infty$ they retain their two-fold degeneracy and nonzero order parameter ($\sum_i S_i^\lambda \geq SN/2$).

The three AIAO_λ phases are stable in a small magnetic field. However, if the respective coupling J_λ is sufficiently weak (i.e. if $|J_\lambda| < J_{\lambda'}$ for some $\lambda' \neq \lambda$), then the order is destroyed at a finite field, marking a transition to a $\text{Chain}_{\lambda'}$ phase. Note that in a finite field, the AIAO_λ states will acquire a finite magnetization since the magnetic moments on the α chains will be partially polarized. However, even as $h \rightarrow \infty$, the phase retains an intensive two-fold degeneracy and a nonzero order parameter $\sum_i S_i^\lambda \geq SN/2$. These two AIAO states are related by a mirror symmetry in the case of the dipolar AIAO_X and AIAO_Z phases and by a combination of time-reversal and a mirror symmetry in the case of the octupolar AIAO_Y phase.

The extensive degeneracy of the three spin ice phases SI_λ at any finite field is broken down to a subextensive degeneracy, yielding effectively one-dimensional degrees of freedom. We first consider the dipolar cases, i.e. $\lambda = x, z$, shown in Fig. 9.1 (b): while the magnetic moments on the α chains (sublattices 0 and 3) are polarized by the field, magnetic moments on the β chains do not couple to the field at all, but are restricted to two possible configurations per β chain by the two-in-two-out rule. Thus, each β chain carries an independent effective Ising degree of freedom. We denote this phase when resulting from dimensional reduction in a SI_λ phase by Chain_λ . In all of them, there is a finite magnetization in the direction of the field, stemming from the polarized α chains and each β chain carries a finite magnetization in the direction of the chain multiplied by its respective Ising degree of freedom.

9.1.3. Octupolar chain phase

The case of the octupolar spin ice phase SI_Y is special since the local moment does not couple to the field. Because of this, the low-field phase which evolves out of the octupolar ice SI_Y is distinct from the other chain phases. If the octupolar exchange J_y dominates, a weak field partially polarizes the magnetic moments on the α chains, but leaves a finite octupolar component S_i^y . This component does not couple to the external field even on the α chains. The “spin-ice” rule $\sum_{\text{tet}} S_i^y = 0$, however, enforces ordering of those moments along both α and β chains. Hence, again the extensive degeneracy is broken down to a subextensive degeneracy, but with an Ising degree of freedom residing on each α and each β chain. We call this phase Chain_Y^* .

The additional Ising degrees of freedom in the Chain_Y^* phase are related to the staggered octupolar moment on the α chains, which on a chain C is given by

$$m_y^{\text{Staggered}} = \frac{1}{2S} \left(\sum_{i \in \mathcal{L}_0 \cap C} S_i^y - \sum_{j \in \mathcal{L}_3 \cap C} S_j^y \right). \quad (9.6)$$

As the field h is increased, or the octupolar exchange coupling J_y is lowered, this quantity vanishes in a second order transition [Fig. 9.3], marking the onset of a chain phase Chain_Y which is equivalent to the dipolar chain phases $\text{Chain}_{X/Z}$ in that there is exactly one Ising degree of freedom per β chain.

Stability with respect to quantum fluctuations

It is also possible to derive the existence of the $\text{Chain}_{X/Z}$ phase in degenerate perturbation theory from the limit of the octupolar ice phase SI_Y . As we show in the following, for $S = 1/2$ there is a diagonal perturbation at fourth order in the external field that generates an effective antiferromagnetic (with respect to the local

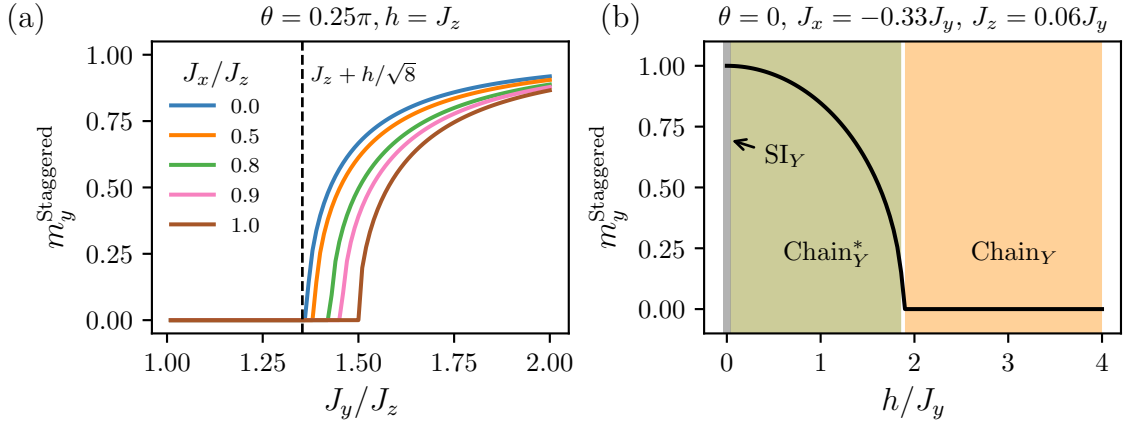


Figure 9.3.: Transition from the Chain_Y to the Chain_Y^* phase. Panel (a) shows it as function of J_y for fixed field $h = J_z$ and $\theta = \pi/4$. Panel (b) shows the same transition as a function of the external field, for exchange parameters used to model $\text{Ce}_2\text{Sn}_2\text{O}_7$ in Ref. [219]. Both plots show the staggered octupolar moment on the α chains [Eq. (9.6)], which is nonzero only in the Chain_Y^* phase. Classically, the octupolar ice phase SI_Y exists only for strictly zero field $h = 0$, however in the quantum case there would be a finite window at low field with a quantum spin liquid ground state.

basis) interaction within the α chains. This interaction favours a set of states with one-dimensional degrees of freedom on each α and β chain, corresponding precisely to the Chain_Y^* phase found in the classical analysis.

For simplicity, we set the transverse coupling to zero and study the Hamiltonian

$$\mathcal{H} = \mathcal{H}_0 + V, \quad (9.7a)$$

$$\mathcal{H}_0 = J_y \sum_{\langle ij \rangle} S_i^y S_j^y, \quad (9.7b)$$

$$V = -h \sum_{i \in \mathcal{L}_0} (S_i^+ + S_i^-) + h \sum_{i \in \mathcal{L}_3} (S_i^+ + S_i^-). \quad (9.7c)$$

The ground states of \mathcal{H}_0 are given by the octupolar spin-ice configurations, where $\sum_{i \in \text{tet}} S_i^y = 0$ on every tetrahedron.

The leading order correction appears at second order and consists of the perturbation V_i acting on the same site twice

$$\Delta E_i^{(2)} = \langle \psi_0 | V_i (E - H_0)^{-1} V_i | \psi_0 \rangle \quad (9.8)$$

$$= \frac{2Sh^2}{-2SJ_y} = -\frac{h^2}{J_y}. \quad (9.9)$$

The above is just a constant correction to the unperturbed ground state energy E_0 , that is it does not discriminate between ground states ψ_0 , leaving the degeneracy unchanged.

This changes at fourth order, where the only nonvanishing contribution consists of application of the perturbation V_i to two neighboring sites. Note that the perturbation acts only on the α chains. Although in the following we assume that $S = 1/2$,




Figure 9.4.: A possible sequence of intermediate states at fourth order degenerate perturbation theory. The full sequence applied here is $V_i V_j V_j V_i$. Tetrahedra on which the pseudospin two-in-two-out rule is violated are indicated in red, with energy penalties ΔE_{tet} indicated next to them. Note that all valid orders of application of the V_i , V_j will have the same set of intermediate energies since the ground state is not allowed as an intermediate state.

the result is qualitatively the same for other values of S . The energy correction is given by

$$\Delta E_i^{(4)} = 4 \times \frac{(2S)^2 h^4}{-(2SJ_y)^2 SJ_y (3 + \sigma_i \sigma_j)} \quad (9.10)$$

$$= \text{const} + \frac{h^4}{2J_y^3} \sigma_i \sigma_j, \quad (9.11)$$

where $\sigma_i = \pm 1$ denotes the configuration of the octupolar moment on site i in the ground state ψ_0 . The above expression now does discriminate between ground states ψ_0 , favoring those states with a staggered configuration of octupolar moments on the α chains.

The ice rule $\sum_{i \in \text{tet}} S_i^y = 0$ then forces all β chains into a staggered configuration. This leaves an Ising degree of freedom on each of the α chains and each of the β chains, reproducing exactly the degeneracy structure of the classical Chain_Y^* phase.

The above result can be easily understood. First, since any valid sequence of intermediate states flips two spins twice, the numerator is given by $(2S)^2 h^4$. One such sequence is shown in Fig. 9.4, with contributions from each tetrahedron to the energy relative to the ground state indicated in red. Since the ground state is not allowed as an intermediate state, there are exactly three other possible sequences which all have the same sequence of intermediate energies. The first and last state have one of the two moments S_i^y or S_j^y flipped, and hence have an energy of $E - E_0 = 2SJ_y$. The second intermediate state has both moments flipped and thus violates the spin-ice rule on three tetrahedra. Two of those include just one of the two sites and hence contribute also a term $2SJ_y$. The tetrahedron that shares the two sites, however, contributes an energy depending on the ground state configuration. If the two flipped sites were initially oriented antiparallel, then the spin ice rule is still fulfilled on this tetrahedron if both are flipped, while in the case of the two sites being parallel, the spin ice rule is now maximally violated. Altogether the energy of the second intermediate state can thus be written as $E - E_0 = SJ_y(3 + \sigma_i \sigma_j)$.

Adding transverse exchange (J_x, J_z) to Eq. (9.7) will allow the appearance of a ring exchange term in perturbation theory which will favour a $U(1)$ QSL [213]. The competition between this ring exchange and the diagonal term $\Delta E_{ij}^{(4)}$ will determine

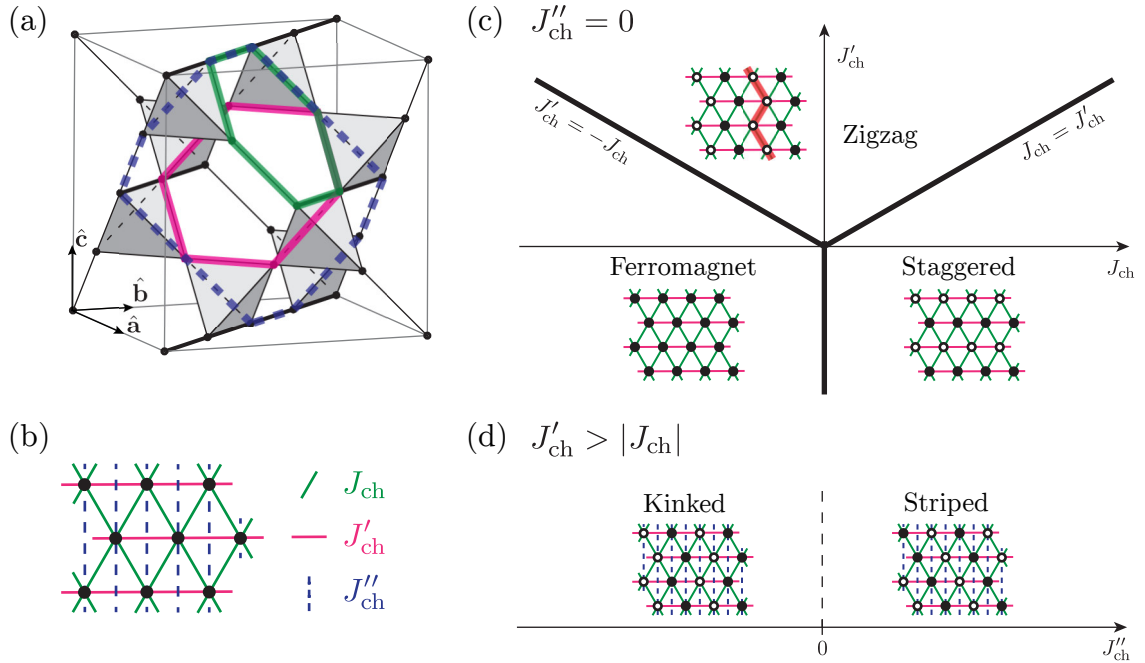


Figure 9.5.: Effective triangular lattice Ising model. (a) The smallest linked clusters that yield a correction to the classical energy and distinguish between classical ground states. The part of the energy correction that distinguishes between classical ground states is proportional to the product of the Ising variables on the chains they connect. (b) The resulting effective anisotropic triangular lattice Ising model. (c) The phase diagram of the anisotropic triangular lattice Ising model as a function of the two nearest-neighbors couplings for the case of negligible next-nearest-neighbor coupling [260]. (d) Lifting of the subextensive degeneracy in the zigzag phase by next-nearest-neighbor coupling J''_{ch} .

the low field phase boundary between the QSL and the Chain_Y^* phase. Since the transition from QSL to Chain_Y^* does not require the condensation of spinons, but instead confines their dynamics, it is a type of confinement transition. This confinement transition should precede the Anderson-Higgs type transition from spinon condensation predicted in gauge mean field theory calculations [247, 259], which generally requires $h \sim J_y$ in order to close the spinon gap.

9.2. Order by Quantum Fluctuations

9.2.1. Effective triangular lattice Ising model

As established in the previous section, if none of the exchange constants J_λ is strongly negative, the system at moderate fields is in a chain phase [Fig. 9.2], where the classical ground state is subextensively degenerate with an Ising degree of freedom carried by each β chain [Fig. 9.1]. The interaction between those Ising degrees of freedom is frustrated and cancels out exactly on the mean-field level. However, since the resulting classical ground state degeneracy is accidental in the sense that the degenerate states are not related by any symmetry of the Hamiltonian [Eq. (9.1)]

one would expect it to be lifted by quantum fluctuations. This is called (quantum) order-by-disorder (OBD) and conceptually it was already discussed in Sec. 2.1.2.

Nearest-neighbor model

In this section, we derive the leading order contribution to the effective interaction between the chains using real space perturbation theory (RSPT) [261–264]. The method divides the Hamiltonian into an unperturbed part, which has a particular classical state as its ground state, and a perturbation incorporating all transverse couplings. One then performs standard perturbation theory to obtain the corrections to the classical energy due to the transverse terms. Since the perturbation depends parametrically on the explicit classical state, the energy correction will depend on it too, possibly lifting classical degeneracies.

We restrict the discussion to the “simple” chain phases Chain_λ where the effective interactions take the form of a nearest-neighbor anisotropic triangular-lattice Ising model.

We can parameterize any classical ground state $\{\mathbf{S}_i^{(0)}\}$ by L^2 Ising variables $\{\eta_i\}$, $\eta_i \in \{+1, -1\}$, where L is the linear system size and $N = 4L^3$ the number of sites in the lattice. We then want to calculate the lowest-order energy correction that distinguishes different ground states, that is, different configurations of $\{\eta_i\}$. As we discuss in detail in App. F, such corrections arise from linked clusters of nontrivial topology on the pyrochlore lattice, the smallest of which are hexagons [Fig. 9.5 (a)]. The leading correction distinguishing configurations of $\{\eta_i\}$ thus occurs at sixth order in RSPT.

Any hexagon on the pyrochlore lattice connects exactly two β chains, but there are two inequivalent kinds of hexagon, as shown in Fig. 9.5 (a). Neglecting constant terms, that is terms that do not distinguish different chain configurations $\{\eta_i\}$, the energy correction from a hexagon is proportional to the product of the Ising variables of the two chains i, j it connects

$$\delta E_{ij}^{(6)} = \text{const} + J_{\text{ch}} \eta_i \eta_j, \quad (9.12a)$$

$$\delta E'_{ij}{}^{(6)} = \text{const} + \frac{J'_{\text{ch}}}{2} \eta_i \eta_j, \quad (9.12b)$$

where we use J_{ch} and J'_{ch} to distinguish the two ways in which two nearest-neighbor chains can be connected by a hexagon. The factor of $1/2$ here is needed since there are $2L$ equivalent hexagons connecting two chains i and j in the primed direction (indicated in pink in Fig. 9.5). The interaction energy between two β chains i, j is then

$$E_{ij}^{(\text{int})} = L J_{\text{ch}}^{(l)} \eta_i \eta_j, \quad (9.13)$$

where $2L$ is the number of sites on the chain.

The β chains embedded in the pyrochlore lattice thus form an anisotropic triangular lattice Ising model which can be described using an effective Hamiltonian

$$\mathcal{H}_{\text{eff}} = J_{\text{ch}} \sum_{\langle ij \rangle} \eta_i \eta_j + J'_{\text{ch}} \sum_{\langle ij \rangle'} \eta_i \eta_j, \quad (9.14)$$

where the sum over $\langle ij \rangle$ runs over nearest neighbor pairs in the $[112]$ and $[\bar{1}\bar{1}2]$ directions and the sum over $\langle ij \rangle'$ indicates summation over nearest-neighbor pairs

in the [110] direction. Note that here, we use “nearest-neighbor” to mean “connected by a hexagon” instead of referring to real-space distance since the β chains are farther apart in the [110] direction than in the other two directions. We indicate which chain coupling is generated by which linked cluster by matching colors in Fig. 9.5 (a) and (b).

The effective couplings $J_{\text{ch}}^{(\prime)}$ are functions of the original pyrochlore exchange couplings J_λ , the field h and the mixing angle θ . The explicit functional form is different in each of the different chain phases. We thus denote the chain couplings in the Chain_λ phase by $J_{\text{ch}}^{\lambda(\prime)}$:

$$J_{\text{ch}}^{(\prime)} = J_{\text{ch}}^{\lambda(\prime)} \quad \text{in } \text{Chain}_\lambda \text{ phase.} \quad (9.15)$$

We calculate the $J_{\text{ch}}^{\lambda(\prime)}$ in sixth-order RSPT (see App. F for details on the perturbative expansion). First consider the result in the Chain_Z phase

$$J_{\text{ch}}^z = -S J_z \cos(\phi)^2 \frac{J_x^2 J_y^2 (J_x - J_y)^2}{J_z^4} \frac{1}{J_z^2} \frac{1}{16} \Gamma_z(h), \quad (9.16a)$$

$$J_{\text{ch}}^{z'} = S J_z \cos(\phi)^4 \frac{J_x^2 J_y^3}{J_z^5} (J_x/J_z + \tan(\phi)^2) \frac{2}{16} \Gamma'_z(h), \quad (9.16b)$$

where S is the spin length and the J_λ are the exchange couplings in Eq. (9.1). The angle ϕ parameterizes the classical configuration of spins on the α chains such that $\theta - \phi$ is the angle between the local field and the magnetic moments. ϕ thus depends on the external field h, θ as well as on the couplings J_x, J_z and $\phi \rightarrow \theta$ as $h \rightarrow \infty$. Γ_z and Γ'_z are dimensionless, positive factors controlling the asymptotic field dependence of the couplings. While full expressions are given in App. F, we note that the field dependence of Γ_z and Γ'_z is qualitatively different:

$$\Gamma_z(h) \sim h^{-2} \quad \text{as } h \rightarrow \infty, \quad (9.17a)$$

$$\Gamma'_z(h) \sim h^{-4} \quad \text{as } h \rightarrow \infty, \quad (9.17b)$$

where the asymptotic exponents corresponds to the number of polarized sites on the hexagon generating the respective coupling [Fig. 9.5 (a)]. This enables tuning the order-by-disorder strength and even the ground state by means of the magnetic field (see Sec. 9.2.3 for details).

The respective expressions for J_{ch}^x and $J_{\text{ch}}^{x'}$ are obtained from Eq. (9.16) by interchanging $J_x \leftrightarrow J_y$ and $\cos \phi \leftrightarrow \sin \phi$. It is also given explicitly in App. F.

In the Chain_Y phase, the couplings in sixth-order RSPT are given by

$$J_{\text{ch}}^y = -\frac{1}{64} S J_y \sin(2\phi)^2 \frac{J_x^2 J_z^2 (J_x - J_z)^2}{J_y^6} \Gamma_y(h), \quad (9.18a)$$

$$J_{\text{ch}}^{y'} = 0, \quad (9.18b)$$

where Γ_y again is a dimensionless, positive, field dependent factor that vanishes as h^{-2} for $h \rightarrow \infty$. For $S = 1/2$, a finite contribution to $J_{\text{ch}}^{y'}$ even at higher orders is only generated when including a transverse term corresponding to cubic magnon terms in linear spin wave theory (LSWT).

The phase diagram of the effective model [Eq. (9.14)] in terms of the effective couplings J_{ch} and J'_{ch} [260] is shown in Fig. 9.5 (c). There are three possible phases: if J_{ch} is negative and $-J_{\text{ch}} > J'_{\text{ch}}$ the system orders ferromagnetically. If the exchange parameter J_{ch} is positive and larger than J'_{ch} , the system enters a ‘staggered’

phase where there is ferromagnetic order along the primed direction (indicated in Fig. 9.5 (b) in pink) with antiferromagnetic order in the perpendicular direction. Finally, if $J'_{\text{ch}} > |J_{\text{ch}}|$, the system is frustrated since when the system is ordered antiferromagnetically in the primed direction (indicated in Fig. 9.5 (b) in pink), it is not possible to satisfy the non-primed bonds (indicated in Fig. 9.5 (b) in green). Thus, in this ‘zigzag’ phase, the ground state is highly degenerate, with each antiferromagnetically ordered row retaining an effective Ising degree of freedom. Following the ferromagnetic order in the non-primed direction yields the eponymous zigzag shape, indicated in red in the sketch in Fig. 9.5 (c).

Since it is clear from Eq. (9.16) and Eq. (9.18) that J_{ch} is negative for all parameters, the chains will never order in a staggered configuration. This already yields a ferromagnetic ground state across the whole Chain_Y phase since there $J'_{\text{ch}} = 0$ for $S = 1/2$ at sixth order. In contrast, in the Chain_X and Chain_Z phases J'_{ch} is finite and can take either sign. The system is thus either in the ferromagnetic or zigzag phase, depending on sign and strength of J'_{ch} relative to J_{ch} . Focusing on the Chain_Z phase for simplicity, in terms of the original pyrochlore exchange couplings (J_λ), if J_y and $J_x + J_z \tan(\phi)^2$ have different signs, J'_{ch} is negative and we expect the chains to order ferromagnetically. In contrast, for $J_x = J_y$, J_{ch} vanishes while J'_{ch} is positive and finite, at least for $J_x > 0$ and $J_x < -J_z \tan^2 \phi$ so we expect the chains to order in a zigzag configuration. As indicated in Fig. 9.5 (c), the phase boundary between the ferromagnetic and zigzag phases is obtained by equating $J'_{\text{ch}} = -J_{\text{ch}}$. This phase boundary as a function of the J_λ is also indicated in Fig. 9.7 as a solid line.

Note that, while the leading order contribution to the effective chain couplings vanishes in certain regions of the phase diagram, we expect that higher-order terms will contribute also in those regions and eventually lift the classical degeneracy. However, since the degeneracy lifting in these cases is of higher order, it is strongly suppressed.

Degeneracy lifting in the zigzag phase

Dipolar-octupolar pyrochlores in a magnetic field undergo a remarkable dimensional evolution when considering different energy scales. On the largest energy scale, corresponding to the bare parameters of the Hamiltonian, the system appears fully three dimensional. In an intermediate regime (between J_{ch} , J'_{ch} and J_z) the system undergoes a dimensional reduction such that it can be described as an ensemble of noninteracting one-dimensional chains. At the lowest energy scales, one might expect the order by disorder mechanism to restore the three dimensional nature of the system.

The zigzag phase, however, in some region of the phase diagram adds an additional stage to this dimensional evolution. Viewed in the pyrochlore lattice, left-over degrees of freedom are effectively two-dimensional. Magnetic moments on the β chains are ordered in stripe patterns within the $[h\bar{l}0]$ plane, but the effective interactions between the planes are frustrated. Thus, the system has a ground state with a subextensive degeneracy scaling with the linear system size L , each plane carrying an effective Ising degree of freedom.

This degeneracy is broken only at 10th order in perturbation theory, by a linked cluster indicated by the blue dashed line in Fig. 9.5 (a). This cluster generates a next-nearest-neighbor coupling J''_{ch} in the $[001]$ direction as indicated in Fig. 9.5 (b)

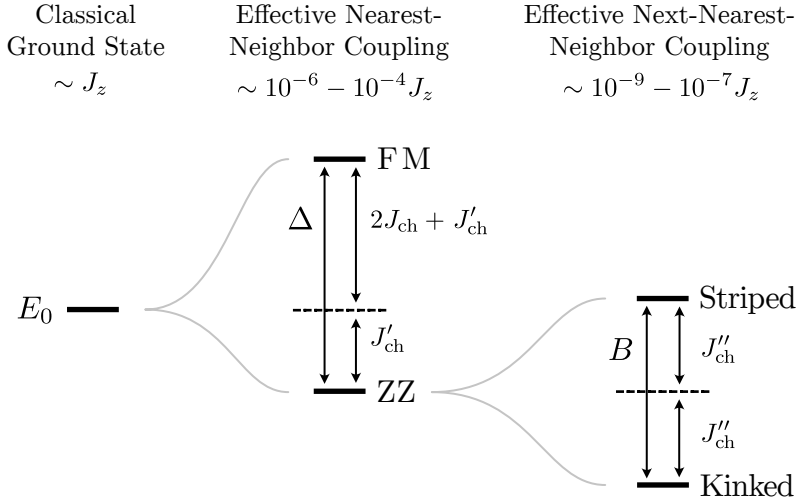


Figure 9.6.: Hierarchy of energy scales in ground state selection in the Chain_z phase. The classical ground state energy is separated from the excited states on a scale set by J_z . On an energy scale of Δ , which for most values of exchange couplings lies in a range of $10^{-6} - 10^{-4} J_z$ when $h \sim J_z$, quantum fluctuations split the classical ground state manifold into the ferromagnetic state and the band of zigzag states. The choice of a single zigzag state happens again on a much lower energy scale B , which for most values of parameters lies in a range of $10^{-9} - 10^{-7} J_z$. The numerical values for the scales are obtained from linear spin wave theory [Fig. 9.7].

and, equivalently, a ring exchange term. Note that the other possible next-nearest-neighbor coupling, while also generated at 10th order, is still frustrated in the zigzag phase and does not lift the degeneracy. As shown in Fig. 9.5 (d), depending on the sign of the next-nearest-neighbor coupling J''_{ch} , there are two different zigzag states selected, either a ‘kinked’ state, where ordered moments form a zigzag pattern in the [001] direction, or a ‘striped’ state, where magnetic moments order ferromagnetically along one of the non-primed nearest-neighbor directions.

The ‘ferromagnetic’, ‘kinked’ and ‘striped’ configurations all have only an intensive number of degenerate states related by symmetries of the Hamiltonian [Eq. (9.1)]. Hence, the system in the zigzag phase regains its full three-dimensional correlations on an energy scale $\sim J''_{\text{ch}}$ while on scales between J''_{ch} and J_{ch} , J'_{ch} it features effectively two-dimensional degrees of freedom.

We summarize the resulting hierarchy of energy scales in ground state selection in Fig. 9.6, where we also give rough numerical ranges in which the scales Δ and B lie for most values of exchange couplings if $h \sim J_z$ [cf. Fig. 9.7].

9.2.2. Order by disorder phase diagram

The classical phase diagram of dipolar-octupolar pyrochlores in a moderate external field in the [110] direction, as discussed in Sec. 9.1, contains extended regions of chain phases in which the field generates effective one-dimensional degrees of freedom and which feature a subextensive degeneracy. Since the degenerate states are not related by symmetries, one expects the degeneracy to be lifted by quantum fluctuations. In most of these chain phases (the Chain_λ phases), this degeneracy lifting, as estab-




Figure 9.7.: Numerical order-by-disorder (OBD) phase diagram of dipolar-octupolar pyrochlores for different mixing angles θ and external fields h , together with the two relevant energy scales Δ and B of ground state selection (cf. Fig. 9.6). We indicate the ground state found by linear spin wave theory (LSWT) by color. ‘Degen’ denotes that all ground states are degenerate, i.e. no OBD is observed. ‘ZZ’ denotes that a zigzag state is selected, but that the zigzag states are degenerate within the numerical precision (cf. the right column in the respective regions). ‘FM’, ‘Str’, and ‘Kink’ denote that ferromagnetic, striped and kinked chain order is selected [cf. Fig. 9.5(d)]. We also indicate the transition between the FM and ZZ phases as computed from real space perturbation theory (RSPT) by a black line.

lished in Sec. 9.2.1, can be cast in the form of an anisotropic triangular-lattice Ising model with three effective chain coupling parameters. Motivated from the effective model, which was derived in leading-order real-space perturbation theory (RSPT), we expect that on the largest energy scale, two different types of chain configuration will be chosen as a ground state. The chains will order either ferromagnetically, or in a zigzag state [Fig. 9.5 (c)]. In the latter case, there is a further dimensional evolution to effectively two-dimensional degrees of freedom, which are planes perpendicular to the [001] direction, each carrying an effective Ising degree of freedom, noninteracting in leading-order RSPT. An effective interaction J''_{ch} , breaking the residual degeneracy, is only generated at 10th order of RSPT. Depending on the sign of J''_{ch} , one of the two zigzag configurations will be chosen, either a ‘kinked’ or a ‘striped’ state both of which then feature an intensive number of degenerate states related by symmetries [Fig. 9.5 (d)].

Phase diagram

In this section, we turn to calculate the full classical degeneracy lifting in the Chain_λ phases, as a function of the exchange parameters J_λ . To corroborate the RSPT calculation, we calculate the lifting numerically using linear spin wave theory (LSWT). The calculation is performed using a semi-infinite slice of a pyrochlore lattice, that is a 8×8 triangular lattice of infinite β chains with periodic boundary conditions in the directions orthogonal to the chains.

As explained in detail in Sec. 2.1.2, the central quantity to compute ground state selection by quantum fluctuations is the zero point energy

$$\mathcal{E}_0(\{\mathbf{S}_i^{(0)}\}) = \frac{1}{\pi} \sum_{\nu} \int_0^{\pi/2} dk \omega_{\nu}((k, -k, 0)), \quad (9.19)$$

where $\omega_{\nu}(\mathbf{q})$ are the magnon frequencies of the system around the ground state $\{\mathbf{S}_i^{(0)}\}$. Details of the numerical computation can be found in App. G. We compute the zero point energy for all possible ground state configurations of the effective nearest-neighbor model, which are the ferromagnetic state, the staggered state and 12 different zigzag states (those are all zigzag states on a 8×8 triangular lattice not related by lattice symmetries). Assuming that the effective model is a valid low-energy description of the system, this will yield the correct ground state phase diagram and serve as a quantitative test of the effective couplings calculated in RSPT. To test the validity of the effective model, we also compute the zero point energies for 1000 randomly selected chain configurations and compare them to the predictions of the effective model, with three couplings fitted using LSWT (see below for details). We find that the effective model indeed gives an accurate estimate of the energies, and therefore that searching for ground states only among the possible ground states of the effective model is justified.

The resulting (quantum) order-by-disorder phase diagram for the Chain_λ ($\lambda = x, y, z$) phases is shown in Fig. 9.7 for different mixing angle $\theta = 0.25\pi$ (a-b) and $\theta = 0.33\pi$ (c-d) as well as for different fields $h = 0.5$ (a, c) and $h = 4$ (b, d). The selected configurations are those expected from the effective Ising model. Remarkably, the phase boundaries obtained from LSWT also match very well with those obtained directly from RSPT at least for moderate transverse couplings. Going beyond the RSPT calculation, LSWT also reveals the splitting of the zigzag state into the

striped and kinked configurations. There are also parts of the phase diagram where the classical degeneracy is not lifted within LSWT. Most notably, this occurs along the two lines $J_x = 0$ and $J_y = 0$ for which we prove explicitly that the dispersion $\omega(\mathbf{q})$ is identical for all classical ground states within LSWT in App. G.

To shed light on the relevant energy scales in the ground-state selection by quantum fluctuations, we also show in Fig. 9.7 the gap between the ferromagnetic and the zigzag configurations Δ as well as the bandwidth of the zigzag states B (that is also the gap between the striped and the kinked state). It is evident that the energy scales on which the classical degeneracy is lifted are very small and have a clear hierarchy as illustrated in Fig. 9.6. In particular, for the compound $\text{Nd}_2\text{Zr}_2\text{O}_7$, using the most recent estimate of the exchange couplings [265] and a field of about 1 Tesla yields $\Delta \approx 10^{-8}$ meV. In this example, quantum OBD has thus no experimental relevance in ground state selection.

To investigate the influence of finite temperature on OBD, we employ two approaches (see Sec. G.3 in App. G for details). First, we compute the leading order contribution of the classical low-temperature expansion of the free energy and show that it is identical for all classical ground states in the Chain_λ phases. Second, we compute the free energy of the noninteracting magnon gas at finite temperature and show that below a crossover temperature T_{co} set by the spin wave gap, the energy scales Δ and B are barely modified from their zero temperature values. For $\text{Nd}_2\text{Zr}_2\text{O}_7$, we estimate $T_{\text{co}} \approx 0.23\text{K}$ (see Fig. G.2).

Our findings are consistent with the fact that disordered chain configurations were observed in neutron scattering experiments on this compound [266] at low temperatures.

Validity of effective model

In order to further assess the validity of the effective model, we compare it to the full spectrum of zero point energies. To this end, we compute the effective chain couplings J_{ch} and J'_{ch} explicitly using the zero point energies obtained from LSWT using

$$\begin{aligned} J_{\text{ch}} &= \frac{1}{4} \left(\mathcal{E}_0^{\text{FM}} - \mathcal{E}_0^{\text{Staggered}} \right), \\ J'_{\text{ch}} &= \frac{1}{2} \left(\mathcal{E}_0^{\text{FM}} + \mathcal{E}_0^{\text{Staggered}} \right), \\ J''_{\text{ch}} &= \frac{1}{2} \left(\mathcal{E}_0^{\text{Kinked}} - \mathcal{E}_0^{\text{Striped}} \right). \end{aligned} \quad (9.20)$$

We then compare the zero point energies $\mathcal{E}_0(\{\mathbf{S}_i^{(0)}\})$ of all 1014 classical ground state configurations that we consider (ferromagnet, staggered, 12 zigzag states and 1000 randomly chosen configurations) to the estimate from the effective nearest-neighbor model using the effective couplings obtained from Eq. 9.20.

Within the Chain_Z and Chain_X phases, the estimates agree remarkably well with the energies computed from LSWT, with residues between the two on the order of the next nearest-neighbor coupling J''_{ch} [see Fig. G.1 in App. G]. Restricting ourselves to the 12 zigzag states, their splitting is modeled by the fitted J''_{ch} up to $10^{-12}J_z$.

In the Chain_Y phase, we expect the effective model to be still valid as derived in RSPT. At high fields, both LSWT and RSPT predict dominance of the ferromagnetic configuration for all exchange couplings in that phase. However, at low fields LSWT




Figure 9.8.: Difference of the zero point energy of the ferromagnetic chain configuration and the zigzag configuration from LSWT as a function of field for fixed exchange parameters. (a) for $J_x = 0.9J_z$, $J_y = 0.45J_z$, it is possible to tune the ground state of the system by means of the external field. (b) for $J_x = J_y = 0.8J_z$, the ground state is always a zigzag configuration, however the OBD energy scale is much larger. The splitting of the zigzag states in both cases is imperceptible on the scale of the plot.

for some part of the phase diagram predicts a zigzag ground state, in contrast with RSPT. The disagreement can be rationalized by considering that LSWT is unable to capture the hard-core constraint of magnons for $S = 1/2$ and this constraint leads the coupling J'_{ch} to vanish at sixth order in RSPT. Meanwhile, in RSPT we neglect higher orders which could lead to a finite J'_{ch} even for $S = 1/2$.

9.2.3. Field tunability

As is qualitatively clear from Eq. (9.17), the effective nearest-neighbor couplings depend on the external magnetic field h . Since the effective interaction between chains is generated by fluctuations on the α chains, the external field can directly control the strength of these fluctuations and hence be used to manipulate the OBD effect. This is clearly seen in Fig. 9.8, where the difference of the zero point energies of the ferromagnetic state and the zigzag states is shown as a function of field for two sets of exchange parameters.

Furthermore, the ground state that is selected by quantum fluctuations can itself be tuned by means of the field. This is the case because selection of the ferromagnetic state or the zigzag states is determined by a competition of the effective couplings in the two inequivalent nearest-neighbor directions of the lattice. Inspecting the two leading-order contributions shown as green and pink hexagons in Fig. 9.5 (a), we see that the two hexagons have two and four sites located on α chains respectively. Hence, as is reflected in the different behaviour of Γ_z and Γ'_z in Eq. (9.17), the couplings scale differently with the field and hence the competition between J_{ch} and J'_{ch} changes as a function of field. As an example of this, we show in Fig. 9.8 (a) how the energy difference between the ferromagnetic state and the zigzag states (computed using LSWT) changes sign as a function of field.

9.3. Experimental Relevance

Before discussing the concrete modeling of $\text{Ce}_2\text{Zr}_2\text{O}_7$ in finite field in the next section, we here discuss the implications of our results for the dipolar-octupolar pyrochlore compounds more generally.

9.3.1. Disordered dipolar chain state in $\text{Nd}_2\text{Zr}_2\text{O}_7$

First, a transition from an AIAO ordered phase to a disordered chain phase induced by a field has been observed in $\text{Nd}_2\text{Zr}_2\text{O}_7$ [266]. Such a transition is expected for $-J_x < J_z < -J_x/3$ or $-J_z < J_x < -J_z/3$, which is consistent with the exchange parameters of this compound as estimated from neutron scattering, for which the most recent estimate (taken from Ref. [265]) is

$$J_x \approx 0.1 \text{ meV}, \quad J_z \approx -J_x/2, \quad J_y \approx 0.15J_x. \quad (9.21)$$

For these exchange parameters and assuming a field of 1 Tesla, LSWT predicts a ferromagnetic order of the chains that is, however, separated from the zigzag state by a gap of only $\Delta \approx 10^{-8}$ meV. At temperatures above $T_{\text{co}} \approx 0.23 \text{ K}$, this is even further suppressed (see discussion in Sec. G.3 of App. G). Hence, even at the lowest experimentally accessible temperatures one would not expect to see ordering of the chains driven by quantum fluctuations, which is consistent with the experimental observation of a disordered chain state.

This makes $\text{Nd}_2\text{Zr}_2\text{O}_7$ in a moderate [110] field an excellent platform for the study of one-dimensional quantum XYZ chains.

9.3.2. Proposed octupolar spin liquids in $\text{Ce}_2\text{Zr}_2\text{O}_7$ and $\text{Ce}_2\text{Sn}_2\text{O}_7$

Second, recent experimental evidence for an octupolar $U(1)$ quantum spin liquid (QSL) has been reported in the dipolar-octupolar pyrochlores $\text{Ce}_2\text{Zr}_2\text{O}_7$ (see Sec. 8.2) as well as in $\text{Ce}_2\text{Sn}_2\text{O}_7$ (Ref. [219]). While our results presented in Sec. 8.3 contradict the claim of Ref [219] that $\text{Ce}_2\text{Sn}_2\text{O}_7$ has dominantly octupolar exchange, it is nevertheless useful to explore what would happen if the scenario in Ref [219] is realized. The case of $\text{Ce}_2\text{Zr}_2\text{O}_7$ will be discussed in more detail in Sec. 9.4.

If a dipolar-octupolar material realizes an octupolar spin liquid, this implies that in this compound octupolar exchange is positive and large [227, 228, 247, 259]

$$J_y > \max(J_x, J_z, -3J_x, -3J_z). \quad (9.22)$$

Classically, this would result in an ice-like phase of octupolar moments SI_Y at zero-field. This phase has the same extensive degeneracy as spin ice. At finite fields, the system would be in the Chain_Y^* phase in which the extensive degeneracy is broken down to a subextensive degeneracy with one Ising degree of freedom carried by each α and each β chain [Fig. 9.1]. At a finite critical field, the system would then undergo a second-order transition to the Chain_Y phase, which also features a subextensive degeneracy but with only half as many Ising degrees of freedom, located on the β chains. This is shown in Fig. 9.3 (b), where we plot the (classical) octupolar staggered moment [Eq. (9.6)] as a function of the external field for exchange parameters as used to model $\text{Ce}_2\text{Sn}_2\text{O}_7$ in Ref. [219]. The second order transition at finite

field is absent if the ground state is a dipolar spin ice state SI_X and SI_Z since those can be directly connected to the respective chain phases Chain_X and Chain_Z .

In a real compound, where $S = 1/2$, quantum fluctuations might be expected to be strong, the story laid out above gets slightly modified. The zero field octupolar ice phase SI_Y will be replaced by a $U(1)$ QSL that prevails for some finite range of external field [196]. At some critical field $h_c^{(1)}$ this is followed by a confinement transition to a quantum version of the Chain_Y^* phase. At a larger, second critical field $h_c^{(2)}$, there will then be a transition to the ‘simple’ chain phase Chain_Y .

The transition from the QSL to Chain_Y^* can be captured by means of degenerate perturbation theory around the low field limit. As was shown in Sec. 9.1.3, the Chain_Y^* also exists in the quantum limit, and should appear in octupolar quantum spin ices subject to moderate [110] fields.

The double transition from a $U(1)$ QSL to the Chain_Y^* phase and finally to the Chain_Y phase as a function of a [110] field appears to be unique to the case of dominantly octupolar exchange and can hence serve as another experimental test for compounds suspected to realize such physics. Additionally, the existence of the Chain_Y^* phase is interesting itself. It serves as an example of how uniform external fields acting on noncollinear magnetic states can lead to the development of complex degrees of freedom. While the number of degrees of freedom grows only with the square of the linear system size L^2 , their interaction should not be expected to be modeled by a local two-dimensional effective model since α and β chains interpenetrate in a nontrivial way.

9.4. Modeling $\text{Ce}_2\text{Zr}_2\text{O}_7$ at Finite Field

In this section, we model the response of the quantum spin ice candidate material $\text{Ce}_2\text{Zr}_2\text{O}_7$ to an external field in the $[1\bar{1}0]$ and $[001]$ directions. As for the zero field case described in Sec. 8.2, we use a combination of numerical linked cluster expansion and semiclassical molecular dynamics to compute static and dynamic properties, respectively. In the high-field limit, we also use Fermionization via the Jordan-Wigner transform to model the effectively one-dimensional chains. In all cases, we confirm the model parameters obtained from the zero-field fit. In particular, we give strong evidence that correlations in $\text{Ce}_2\text{Zr}_2\text{O}_7$ are dominantly octupolar and that the transverse exchange parameters have small anisotropy $|J_x - J_y| \lesssim 0.02$ meV.

The following people contributed to the experimental data shown in this section: **Sample Preparation:** J. Dudemaine, E. M. Smith, J. Gaudet, A. Fitterman, J. Beare, G. M. Luke, B. D. Gaulin, and A. D. Bianchi. **Neutron scattering:** E.M. Smith, J. Gaudet, C. R. C. Buhariwalla, A. Podlesnyak, Guanhyoung Xu, and B. D. Gaulin. **Specific heat:** D. R. Yahne, T. DeLazzer, E. M. Smith, R. Movshovich, B. D. Gaulin, and K. A. Ross

9.4.1. (Re-)Fitting specific heat at finite field

The specific heat capacity of $\text{Ce}_2\text{Zr}_2\text{O}_7$ as a function of external field aligned with the $[1\bar{1}0]$ crystallographic direction is shown in Fig. 9.9. As the external field strength is increased, the specific heat peak moves to higher temperatures and for the largest field value $|\mathbf{H}| = 2.0$ T a clear shoulder develops at low temperatures. Such a




Figure 9.9.: Specific heat of $\text{Ce}_2\text{Zr}_2\text{O}_7$ in $[1\bar{1}0]$ field (note that the $[1\bar{1}0]$ and $[110]$ directions are equivalent). Markers show experimental results, solid lines an NLCE calculation with the best-fit parameters $(J_x, J_y, J_z) = (0.063, 0.062, 0.011)$ meV. Although the parameters are obtained independently here, they are almost identical to those obtained in Sec. 8.2. Also shown is the specific heat of $\text{La}_2\text{Zr}_2\text{O}_7$, which has the same lattice structure but is non-magnetic.

double-featured structure is indeed expected in a $[1\bar{1}0]$ field. At high temperature, the system is in a paramagnetic state. As the temperature is lowered, the α -chains become polarized and align with the external field at a temperature $T \sim g_z \mu_B h / k_B$. The β -chains do not couple to the external field, and so develop correlations at a field independent temperature controlled by the exchange interactions. We thus interpret the double-featured structure as evidence for field-induced dimensional reduction, with the first and second peak corresponding to the onset of correlations on the α - and β -chains, respectively.

We also refit the specific heat using a procedure analogous to that used for the *magnetic susceptibility* in Sec. 8.3. This is necessary since at finite field $|\mathbf{H}| > 0$, the specific heat does depend on the form of the local magnetic moment $\boldsymbol{\mu}_j$ and hence on the value of the mixing angle θ as well the permutation of exchange parameters. In total we perform four full parameter sweeps at the four field values $|\mathbf{H}| = 0, 0.5, 1.0, 2.0$ T. We use sixth-order NLCE for the parameter sweep at zero field and fifth-order NLCE for the sweeps at finite field. The reduced order at finite field is owed to a combination of the additional θ dimension of the parameter sweep and the reduced symmetry of the Hamiltonian, which increased the computation cost of the exact diagonalization step [253]. As in chapter 8, the NLCE calculations were performed by Robin Schäfer. The best-NLCE-fit parameters

$$(J_x, J_y, J_z) = (0.063, 0.062, 0.011) \text{ meV} \quad \theta \approx 0 \quad (9.23)$$

are extremely close to the parameters obtained from fitting the zero field data in Sec. 8.2.

The specific heat as computed from NLCE using the best-fit parameters [Eq. (9.23)] is also shown in Fig. 9.9 as solid lines. For the zero-field calculation, we show both a sixth- and seventh-order calculation, which yields an estimate of its convergence. While the calculation captures the high-temperature behavior above $T > 0.5$ K very well, there is clear disagreement between calculation and measurement below 0.5 K. This is the case also at large fields, most notably for $|\mathbf{H}| = 2$ T, where NLCE predicts the onset of a much larger second peak. Since the fit parameters are the best fit, these disagreements clearly point towards the relevance of effects beyond the nearest-neighbor exchange Hamiltonian in Eq. (9.1). Possible explanation include interactions beyond nearest-neighbor [226] and disorder, a detailed investigation of these possibilities, however, is beyond the scope of this work.

The fit to the in-field heat capacity presented here gives an estimate of the permutation of exchange parameters and the value of θ that is independent from the earlier fits to the magnetic susceptibility and the inelastic structure factor which were presented in Sec. 8.2. Leaving aside questions regarding the effects of disorder and further-neighbor interactions, it thus constitutes independent evidence that the largest exchange couplings correspond to interactions between octupolar moments. Note that similar to what we observed when fitting the zero field data (Sec. 8.2), the specific heat fit does not strongly constrain the exact value of the anisotropy between J_x and J_y and mixing angle. More concretely, increasing θ while at the same time increasing J_y and reducing J_x preserves the octupolar nature of the correlations and hence produces a similar response. As we will see in the following, modeling of the structure factor allows for constraining this anisotropy more quantitatively.

9.4.2. Evidence for the Chain_Y^* phase?

As discussed in Sec. 9.1.3, the evolution from spin ice to the high-field chain state happens very distinctively in the case of octupolar spin ice. At low field, we predict the existence of an intermediate phase where already spinon dynamics remains only along, rather than between, chains but both the unpolarized β chains as well as the polarized α chains carry an effective Ising degree of freedom. This phase, which we call Chain_Y^* , at larger fields gives way to the usual chain phase at high fields via a gap-closing. The high-field Chain_Y phase is again gapped, but with the ground state of the polarized α chains being unique.

Since we predict $\text{Ce}_2\text{Zr}_2\text{O}_7$ to realize an octupolar spin liquid at zero field, it represents an ideal testbed for the this prediction. Observing the Chain_Y^* is, however, complicated by the fact that also $J_x \approx J_y$, as we will show in the following.

Critical field

We assume that the system at some low field can already be described as two independent sets of chains. In this limit, with $\theta = 0$, we can solve the thermodynamics as well as the dynamics exactly.

To simplify the treatment, we will neglect the weakest exchange parameter, $J_z \approx 0$. In this case, the Hamiltonian on both α - and β -chains reduces to the XY model in a staggered field

$$H = \sum_j J_x S_j^x S_{j+1}^x + J_y S_j^y S_{j+1}^y + (-1)^j h S_j^z. \quad (9.24)$$

with $h = 0$ for the β chains and $h = |\mathbf{H}|\sqrt{6}/(2\mu_B g_z)$ for the α chains as before [Eq. (9.4)]. Here, $|\mathbf{H}|$ is the experimentally applied field-strength, μ_B is the Bohr Magneton and $\sqrt{6}/2$ is a geometrical factor arising from the projection of the external field on the local easy-axis.

Rotating the basis on every second site by π around the local x -axis yields $S^y \rightarrow -S^{\bar{y}}$ and $S^z \rightarrow -S^{\bar{z}}$ on these sites, transforming the above into a chain in uniform field but with flipped sign of the J_y exchange term

$$H = \sum_j J_x S_j^x S_{j+1}^x - J_y S_j^{\bar{y}} S_{j+1}^{\bar{y}} + h S_j^{\bar{z}}. \quad (9.25)$$

This Hamiltonian can be solved straightforwardly by using the Jordan-Wigner transformation and subsequent Bogolyubov transformation [267]

$$H = \sum_k \omega(k) \left(\eta_k^\dagger \eta_k + \frac{1}{2} \right) \quad (9.26)$$

$$\omega(k) = \sqrt{(\gamma \cos(k) + h)^2 + J^2 \sin(k)^2} \quad (9.27)$$

where the η_k are Fermionic operators, $J = \frac{1}{2}(J_x + J_y)$, and $\gamma = \frac{1}{2}(J_x - J_y)$. Note that J and γ switched places in the expression for the dispersion ω compared to the usual result, because of the flipped sign in front of J_y in Eq. (9.25). At zero field $h = 0$, the fermion dispersion is gapped with gap $\Delta = \gamma$ and the ground state of Eq. (9.25) is two-fold degenerate. As the field h is increased, the gap is reduced

until it closes at the critical field $h_c = \gamma$. At $h > h_c$ the dispersion is again gapped but the ground state of Eq. (9.25) is now unique. The critical field $h_c = \gamma$ hence marks the transition from the Chain_Y^* to the Chain_Y phase.

This means that in $\text{Ce}_2\text{Zr}_2\text{O}_7$, for which we estimated $2\gamma = J_x - J_y \approx 0$, the critical field is very low. This complicates the observation of the Chain_Y^* phase but also serves as an independent measure of the anisotropy.

Probing the β -chain magnetization in experiment

A direct signature of the different chain phases in principle is the staggered magnetization of the chains. While the octupolar moment m_{st}^y of the α chain as shown in Fig. 9.3 is not directly accessible to experiment, the transition is also visible in the dipolar moment m_{st}^z . In the rotated basis of Eq. (9.25), this is just the magnetization. This is a four-Fermion operator and can be computed exactly

$$\langle m_{\text{st}}^z \rangle = \langle m^z \rangle \quad (9.28)$$

$$= \frac{1}{2\pi} \int_0^\pi \tanh\left(\frac{1}{2}\beta\omega(k)\right) \frac{\gamma \cos(k) + h}{\omega(k)} dk. \quad (9.29)$$

It can be probed in experiment by taking the square root of the Bragg peak intensity at the [220] position

$$m_{\text{st}}^z \propto \sqrt{S(\mathbf{q} = [220])} \quad (9.30)$$

This is shown in Fig. 9.10. In panel (a), we show the peak intensity at the [220] position as a function of field for different temperatures as measured in elastic neutron scattering on $\text{Ce}_2\text{Zr}_2\text{O}_7$. At the lowest temperature we also compare to the calculation in Eq. (9.29) as well as to the result from semiclassical molecular dynamics simulations (see Appendix E for details). We use the parameters in Eq. (9.23), with $J_z = 0$ for the Jordan-Wigner result. Note that the semiclassical simulation only agrees with the experimental measurements if temperature is significantly renormalized, while the Jordan-Wigner calculation captures the trend much better without such renormalization. We interpret this as evidence for strong quantum fluctuations. Indeed, the evolution of the [020] peak in an external field aligned with the [001] crystallographic direction is shown in the inset of Fig. 9.10 (a). In this case, no renormalization of temperature is needed for the MD simulation, which is shown as an orange-dashed line. In panel (b) and (c) of the same figure, we show Eq. (9.29) and its derivative with respect to field, respectively, for $J = 0.064$ meV and different anisotropies γ . The data is clearly not compatible with a large anisotropy $\gamma > 0.02$ meV and is best described by $\gamma \approx 0.01$ meV. Since at very low fields we expect the description of the system as a set of independent chains to break, however, we take this as evidence only for an upper bound on the anisotropy.

9.4.3. Integrated structure factor

In Fig. 9.11, we show momentum resolved quasielastic neutron scattering data in the $[hhl]$ plane and in an external field aligned with the $[1\bar{1}0]$ direction. Panels (a-c) show experimental data at $T = 0.03$ meV obtained from inelastic measurements and integrated over the energy range -0.2 meV $\leq E \leq 0.2$ meV as well as a small momentum range in the out-of-plane direction $[K\bar{K}0]$, $-0.3 \leq K \leq 0.3$. In all cases, a zero-field dataset at the same temperature was subtracted as background.

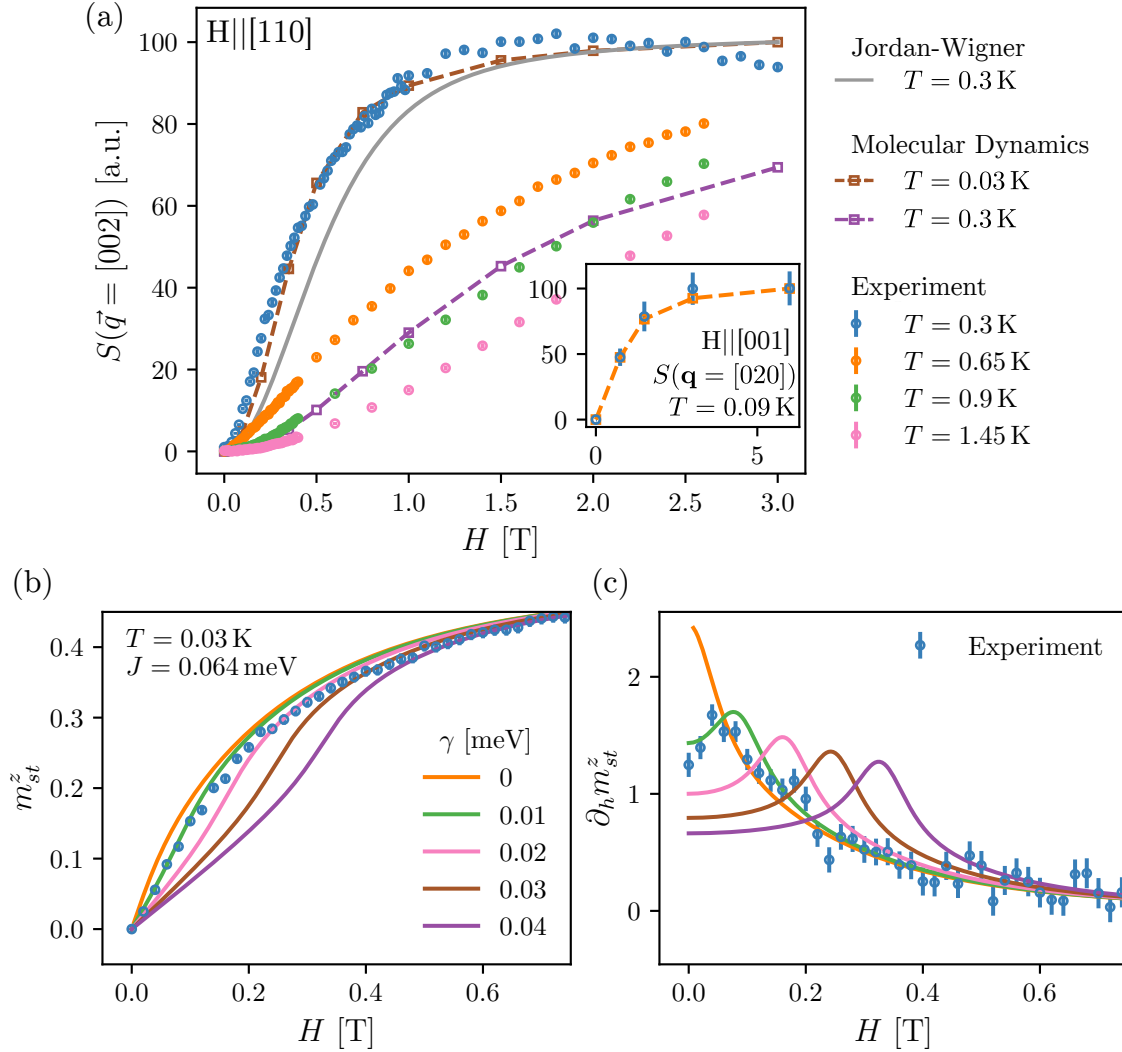


Figure 9.10.: Polarization of $\text{Ce}_2\text{Zr}_2\text{O}_7$ in $[1\bar{1}0]$ field. Panel (a) shows the evolution of the intensity of the magnetic Bragg at $\mathbf{q} = [002]$, measuring the polarization of the α -chains. It is well modelled by both Jordan-Wigner as well semiclassical molecular dynamics (MD) calculations (see main text for details). A significant renormalization of temperature is necessary for the MD. The inset shows the polarization of the sample in a $[001]$ field, where no such renormalization is necessary. Panel (b) and (c) show the staggered pseudospin magnetization of the α chains and its derivative with respect to field, respectively. Markers show experimental results, solid lines show Jordan-Wigner calculations with different anisotropies $\gamma = (J_x - J_y)/2$.




Figure 9.11.: Evolution of the energy-integrated structure factor $\int S(\omega, \mathbf{q})d\omega$ in a $[1\bar{1}0]$ field. Panels (a-c) show experimental results, while panels (d-f) and panels (g-i) show the results of molecular dynamics simulations with mixing angle $\theta = 0$ and $\theta = 0.1\pi$, respectively. The exchange parameters used are those from the best-NLCE-fit.




Figure 9.12.: Energy integrated structure factor $\int S(\omega, \mathbf{q})d\omega$ in the infinite field $h \rightarrow \infty$ limit as computed using the Jordan-Wigner transformation. The exchange parameters used are those from the best-NLCE-fit, but with $J_z = 0$. Only the diffuse β -chain contribution is shown. Remarkably, the result closely resembles that of the semiclassical molecular dynamics simulation, even though the β -chains are close to criticality.

We compare these measurements with the integrated structure factor $S(\omega, \mathbf{q})$ obtained from semiclassical molecular dynamics simulations at the same temperature and integrated over the same energy range but without any integration over out-of-plane momentum and without any background subtraction. The calculations use the best-NLCE-fit parameters in Eq. (9.23) with $\theta = 0$ in panel (d-f) and $\theta = 0.1\pi$ in panel (g-i).

The $\theta = 0$ calculation agrees with the data reasonably well and in particular reproduces the rather broad, rod-like feature along the $[00L]$ direction, as well as the evolution as a function of field strength. The calculation with $\theta > 0$ in contrast shows a sharp line along $[00L]$. This feature is the Bragg-plane associated with scattering from the β chains, which are ordered in one dimension (within the chains) but disordered in the two perpendicular directions (across chains).

As in the discussion in Sec. 8.2, the case $\theta = 0.1\pi$ demonstrates some generic behavior of the system as it develops dipolar correlations. In particular, the central rod-like feature along the $[00L]$ direction is a direct signature of dipolar correlations within the beta chains, its width being inversely proportional to the dipolar correlation length. If correlations are strongly dipolar one expects a sharp, high-intensity rod which however is not observed in Fig. 9.11 (a-c). Such sharp rods have been seen for example in the classical spin ices $\text{Dy}_2\text{Ti}_2\text{O}_7$ [268, 269] and $\text{Ho}_2\text{Ti}_2\text{O}_7$ [269, 270], and also in the dipolar-octupolar pyrochlore $\text{Nd}_2\text{Zr}_2\text{O}_7$ [266], which has a similar size of the magnetic moment as $\text{Ce}_2\text{Zr}_2\text{O}_7$ [265, 266]. Here, instead we see a broad feature of low intensity, suggesting a small correlations length between dipolar moments.

According to our fits to both the heat capacity and chain magnetization, the β chains have small exchange anisotropy and are close to criticality, with power-law




Figure 9.13.: Magnetic field dependence of the powder-averaged neutron inelastic neutron scattering cross section. The data was taken on a single crystal of $\text{Ce}_2\text{Zr}_2\text{O}_7$ aligned with the field along the $[1\bar{1}0]$ (a-c) and $[001]$ (d-f) direction at $T = 0.03$ K and $T = 0.09$ K, respectively. In each case, a data set measured at $h = 0$ T has been subtracted as background.

correlations and strong quantum fluctuations [267]. To check that our semiclassical molecular dynamics simulation still give a reasonably accurate description, we compute also the infinite-field limit of the integrated structure factor using the Jordan-Wigner transform.

To this end, we again approximate $J_z = 0$ and treat correlations within the chains exactly while approximating correlations between chains as the product of expectation values $\langle S_i^\alpha S_j^\beta \rangle = \langle S_i^\alpha \rangle \langle S_j^\beta \rangle$. We then follow the numerical procedure in Ref. [271] to compute dynamical correlations within a single β chain of length $L = 1000$. Incorporating the necessary geometrical factors then yields the full dynamical structure factor $S(\mathbf{q}, \omega)$ of the 3D material in the infinite field limit $H \rightarrow \infty$.

In Fig. 9.12, we show the β chain contribution to the integrated dynamical structure factor for both $\theta = 0$ and $\theta = \pi/10$ (note that exact solution for finite θ is possible only for the β chains). Remarkably, the result is extremely similar to that of the semiclassical Monte-Carlo molecular dynamics simulation, even though the chain is close to criticality at $J_x \approx J_y$.

9.4.4. Dynamics

Finally, we show inelastic (that is energy-resolved) neutron scattering data for $\text{Ce}_2\text{Zr}_2\text{O}_7$ in an external field in Fig. 9.13. The top and bottom row show data for an external field aligned with the $[1\bar{1}0]$ and $[001]$ crystallographic direction, respectively. The signal is clearly resolvable as being out of the elastic line only for

large fields $H \geq 1.5 \text{ T}$, which is consistent with the extremely small estimated energy scales in the exchange Hamiltonian [Eq. (9.23)]. It is also extremely weak and, therefore, clearly visible only after powder-averaging (that is averaging over all momentum directions) the single-crystal data.

Current experiments are not able to resolve the individual dispersion of quasiparticles on the α and β chains. To discuss their dynamics we, instead, resort to semiclassical molecular dynamics simulations and show in Fig. 9.14 and Fig. 9.15 the calculated dynamical structure factor $S(E, \mathbf{q})$ for $\text{Ce}_2\text{Zr}_2\text{O}_7$ in an external field aligned with the $[1\bar{1}0]$ and $[001]$ crystallographic direction, respectively. The simulations use the best-fit parameters in Eq. (9.23) and we show two values of the mixing angle θ . The left column shows the case of dominantly octupolar correlations ($\theta = 0$) and the right column the case of weak dipolar correlations ($\theta = 0.1$).

Starting with the external field aligned with the $[1\bar{1}0]$ direction, in this case we expect the dynamics within the molecular dynamics simulation to separate mostly into one set of magnons confined within the α chains and another set confined within the β chains.

Indeed, this separation is clearly visible as a function of magnetic field. At low field, there is one broad continuum of excitation of low energy and for $\theta = 0.1\pi$, a sharp dispersion is visible on top this continuum. This is because for $\theta = 0$ the correlations are mostly octupolar and hence the dominant contribution to neutron scattering is a two-magnon continuum.

Introducing a weak dipolar component by setting $\theta = 0.1\pi$ makes part of the sharp magnon dispersion, which is present at finite field, visible. As the field is increased, the part of the continuum corresponding to the two-particle continuum of magnons on the α chains moves upwards in energy since the single-particle gap is set by the field. At low energy, a sharp dispersion becomes visible for both $\theta = 0$ and $\theta = 0.1\pi$. This signal corresponds to mixture of scattering off magnons on the β chains.

Note that although the qualitative behavior of quasiparticles is captured by the semiclassical simulation, the dispersions are unlikely to agree very quantitatively with those of the actual quantum system. This is because the particles dispersion is effectively one-dimensional and are hence better described as fractionalized spinon excitations rather than simple magnons [267].

Moving on the external field aligned with the $[001]$ direction, in this case we expect the semiclassical molecular dynamics to give correct quantitative predictions.

The signal in Fig. 9.15 is again separated into multiple contributions. At low energy there is scattering off thermally excited single magnons. This contribution does not move up systematically in energy as the strength of the external field is increased, since it only depends on energy differences within the magnon bands rather than between the magnon bands and the ground state. At finite energy, there is scattering off the two-magnon continuum as well as the sharp single magnon dispersion, both with an energy dependence set by the external field. Remarkably, the single-magnon dispersion is visible also for $\theta = 0$ although in this case the correlations are mostly octupolar. This contribution is from neutrons scattering off the octupole moment carried by the single magnon. Its intensity has a momentum dependence which is distinct from that of the dipolar scattering seen for $\theta = 0.1\pi$. This dependence is set by the octupole contribution to the form factor. For the parameters used here, this octupolar contribution at large fields is stronger than the




Figure 9.14.: Dynamical Structure factor $S(E, q)$ obtained from semiclassical molecular dynamics simulations for the best-fit parameters of $\text{Ce}_2\text{Zr}_2\text{O}_7$ [Eq. (9.23)] at $T = 0.03\text{ K}$ in an external field aligned with the $[1\bar{1}0]$ crystallographic direction. Each row shows different field strength H and the columns show different mixing angles θ .

dipolar scattering off the two-magnon continuum. The presence of significant octupolar scattering at comparatively low momentum transfer is somewhat surprising. While the energy resolution of current experiments does not allow the observation of such a magnon dispersion, it nevertheless opens up another possibility to further investigate $\text{Ce}_2\text{Zr}_2\text{O}_7$ and its sister compounds in the future.

9.5. Conclusion and Outlook

In conclusion, we have mapped out the classical phase diagram of dipolar-octupolar pyrochlores in a $[110]$ field in detail and studied quantum order-by-disorder in the ‘chain’ phases, in which the classical ground state degeneracy can be parametrized by a subextensive number of effective Ising degrees of freedom, each describing the magnetic moment of a one-dimensional chain. We have focused our study on a subset of these phases, where all degrees of freedom are carried by a set of parallel chains. We show that OBD in this case can be modeled remarkably well by a simple effective anisotropic triangular lattice Ising model that we derive in real space perturbation




Figure 9.15.: Dynamical Structure factor $S(E, q)$ obtained from semiclassical molecular dynamics simulations for the best-fit parameters of $\text{Ce}_2\text{Zr}_2\text{O}_7$ [Eq. (9.23)] at $T = 0.09$ K in an external field aligned with the [001] crystallographic direction. Each row shows different field strength H and the columns show different mixing angles θ .

theory and corroborate using numerical linear spin wave theory.

Ground state selection by quantum fluctuations in dipolar-octupolar pyrochlores is a multi-step process governed by a hierarchy of energy scales. First, competition of two inequivalent nearest-neighbor chain interactions drives selection of either a ferromagnetic ordering of the chains or a ‘zigzag’ order. In the latter case, the chains are ordered antiferromagnetically in one direction while the nearest-neighbor coupling in the other two directions is frustrated, leaving a subextensive degeneracy scaling with the linear system size. This left-over degeneracy is finally broken by an effective next-nearest-neighbor coupling, albeit on a minuscule energy scale. Taken together, the system undergoes a highly nontrivial dimensional evolution when considering effective degrees of freedom on different energy scales. Comparing the results of the effective model with LSWT, the two methods mostly agree remarkably well even quantitatively, with phase boundaries between ferromagnetic and zigzag phases almost indiscernible in the dipolar chain phases Chain_X , Chain_Z for moderate transverse couplings.

Highlighting connection to different experiments, we have also discussed that the energy scales on which quantum fluctuations break the classical degeneracies are minuscule for parameter values relevant to $\text{Nd}_2\text{Zr}_2\text{O}_7$. This is consistent with the experimental observation of a disordered chain state in this compound [266] and makes this material a promising platform to study pure quantum XYZ chains. We also study the field-breakdown of the octupolar quantum spin liquid (QSL) state in a $[\bar{1}10]$ field. Generally, we predict that for a field in the $[110]$ direction, the transition to the large-field chain state Chain_Y happens in two steps, via an intermediate phase that we call Chain_Y^* . This sequence of transitions is unique for the case of dominantly octupolar exchange and can hence be used as a diagnostic of these phases. Indeed, modeling new in-field data for the specific heat and structure factor of the material $\text{Ce}_2\text{Zr}_2\text{O}_7$ we find evidence for the existence of the Chain_Y^* phase. Definite observation however is rendered difficult because of the small exchange anisotropy ($J_x \approx J_y$) of the compound. This suppresses the spinon gap and, hence, the critical field at which the transition to the Chain_Y phase occurs. Finally, we predict that high-resolution single crystal neutron scattering might enable the observation of *octupolar* spin wave scattering in finite field. This would allow for precise fitting of exchange parameters but remains out of reach in current experiments.

Appendices

A. Compression protocols for hardcore spins

A.1. Lubachevsky-Stillinger Compression Protocol

Algorithm 1: Lubachevsky-Stillinger compression

Input : Trial increment δ_0
Minimum increment δ_{\min}
Number of rattling sweeps N_{rattle}
Output: Valid configuration $\{\mathcal{S}_i\}$ at Δ_J ,
Final exclusion angle Δ_J

```
1  $\mathcal{S} \leftarrow$  random initial configuration;  
2  $\Delta \leftarrow 0$ ;  
3 while  $\delta > \delta_{\min}$  do  
4    $\delta \leftarrow \delta_0$ ;  
5   while  $\mathcal{S}$  is not valid at  $\Delta + \delta$  do  
6     Perform  $N_{\text{rattle}}$  Monte-Carlo sweeps;  
7      $\delta \leftarrow \delta/2$ ;  
8   end  
9    $\Delta \leftarrow \Delta + \delta$   
10 end  
11 return  $\mathcal{S}, \Delta$ ;
```

The Lubachevsky-Stillinger protocol [63] used in Sec. 3.2.1, was implemented in C++. It is given in pseudocode in algorithm 1.

The fundamental limitation of this protocol is that what determines the largest possible increment of the inclusion angle Δ is the smallest angle between *any* nearest neighbor pair in the system. This value however vanishes in the limit of infinite system size $N \rightarrow \infty$, implying that for fixed N_{rattle} also the effective compression speed $\langle \delta/N_{\text{rattle}} \rangle \rightarrow 0$ as $N \rightarrow \infty$. For $L = 40$ we find $\langle \delta \rangle \approx 10^{-5}$.

A.2. Softcore Protocol

The softcore protocol, originally devised for hard spheres by Xu et al. [65] was implemented in C++ as well. The main idea is to use the gradient of the potential to push overlapping spheres away from each other to produce a valid hard sphere state from a state with local violations of the constraint. The authors then use alternating compression and expansion to produce states where spheres are in close contact with each other.

Algorithm 2: Softcore Compression Step Compress

Input : Valid configuration $\mathcal{S}_0 = \{\mathcal{S}_i\}$ at Δ_0 ,
Trial increment δ_0 ,
Tolerance ϵ

Output: Valid configuration $\{\mathcal{S}_i\}$ at Δ ,
New exclusion angle Δ

```
1  $\mathcal{S} \leftarrow \mathcal{S}_0$ ;  
2  $V \leftarrow V^{(\text{tot})}(\mathcal{S}, \Delta + \delta)$ ;  
3  $\text{MODE} \leftarrow \text{COMPRESS}$ ;  
4  $\delta \leftarrow \delta_0$ ;  
5  $\Delta \leftarrow \Delta_0 + \delta$ ;  
6 while  $V > \epsilon$  or  $V < \epsilon/2$  do  
7    $\mathcal{S}_{\text{old}} \leftarrow \mathcal{S}$ ;  
8    $\mathcal{S}, V, \leftarrow$  Conjugate gradient descent of  $V^{(\text{tot})}(\mathcal{S}, \Delta)$  with initial state  $\mathcal{S}$ ;  
9   if  $V > \epsilon$  then  
10    if  $\text{MODE} = \text{COMPRESS}$  then  
11       $\delta \leftarrow \delta/2$ ;  
12       $\text{MODE} \leftarrow \text{EXPAND}$ ;  
13    end  
14     $\mathcal{S} \leftarrow \mathcal{S}_{\text{old}}$ ;  
15     $\Delta \leftarrow \Delta - \delta$ ;  
16  end  
17  if  $V < \epsilon$  then  
18    if  $\text{MODE} = \text{EXPAND}$  then  
19       $\delta \leftarrow \delta/2$ ;  
20       $\text{MODE} \leftarrow \text{COMPRESS}$ ;  
21    end  
22     $\Delta \leftarrow \Delta + \delta$ ;  
23  end  
24 end  
25 return  $\mathcal{S}, \Delta$  ;
```

Algorithm 3: Softcore Compression Protocol

Input : Max, min step size $\delta_{\max}, \delta_{\min}$
Expansion rattling sweeps $N_{\text{rattle}}^{(\text{exp})}$
Extra rattling sweeps N_{rattle}
Maximum number of expansions $N_{\text{exp}}^{(\text{max})}$
Expansion step δ_{exp}

Output: Valid configuration $\mathcal{S} = \{\mathcal{S}_i\}$ at Δ
Final exclusion angle Δ

```

1  $\Delta \leftarrow 0$ ;
2  $\mathcal{S} \leftarrow$  random configuration;
3 while  $N_{\text{exp}} < N_{\text{exp}}^{(\text{max})}$  do
4    $N_{\text{exp}} \leftarrow 0$ ;
5    $\Delta_{\text{max}} \leftarrow \Delta$ ;
6    $\delta \leftarrow 0$ ;
7   while  $\delta < \delta_{\min}$  and  $N_{\text{exp}} < N_{\text{exp}}^{(\text{max})}$  do
8      $\mathcal{S}, \Delta \leftarrow$  Compress( $\mathcal{S}, \delta_{\max}$ );
9      $\delta \leftarrow \Delta - \Delta_{\text{max}}$ ;
10    if  $\delta < \delta_{\min}$  then
11       $N_{\text{exp}} \leftarrow N_{\text{exp}} + 1$ ;
12       $\Delta \leftarrow \Delta - N_{\text{exp}} \cdot \delta_{\text{exp}}$ ;
13      Perform  $N_{\text{rattle}}^{(\text{exp})}$  Monte-Carlo sweeps;
14    end
15  end
16  Perform  $N_{\text{rattle}}$  Monte-Carlo sweeps;
17 end

```

We implement a compression step closely analogous to [65] for our hardcore spin model as follows, given as pseudocode in algorithm 2. Consider a valid initial state $\mathcal{S} = \{\mathcal{S}_i\}$ at some initial exclusion angle Δ , that is $V^{(\text{tot})}(\mathcal{S}, \Delta) = \sum_{\langle ij \rangle} V(\phi_{ij}, \Delta) < \epsilon$ for some fixed tolerance ϵ , which we always take to be $\epsilon = 2 \times 10^{-10}$. Now, starting with compression, increase the inclusion angle Δ by some trial value $\delta > 0$, and use conjugate gradient descent to get from the state \mathcal{S} , which has now a finite potential energy $V^{(\text{tot})}(\mathcal{S}, \Delta)$, to a final state \mathcal{S}' . Depending on the value of $V^{(\text{tot})}(\mathcal{S}', \Delta)$, we continue with either compression or expansion of the system. If $V^{(\text{tot})}(\mathcal{S}', \Delta) > \epsilon$, we revert the state of the simulation to its state before the energy minimization and reduce the exclusion angle Δ by δ (expansion). If $V^{(\text{tot})}(\mathcal{S}', \Delta) < \epsilon$, we increase the exclusion angle Δ by δ (compression). In both cases, we then minimize the potential energy again by conjugate gradient descent and repeat the procedure. Every time we switch from compression to expansion or vice versa, the increment δ is halved. The step terminates if $\epsilon/2 < V^{(\text{tot})}(\mathcal{S}, \Delta) < \epsilon$. Note that although we use the softened constraint in principle, forcing the system to have zero potential energy is equivalent to imposing the hardcore constraint. Allowing the potential energy to be small ($\epsilon \ll N^{-1}$) but nonzero then ensures that the hardcore constraint is almost fulfilled by the state.

Similar to the LS protocol (algorithm 1), we augment the compression step with local Monte-Carlo dynamics as follows. If the difference between successive exclusion angles Δ is below some threshold δ_{min} , we reduce the exclusion angle Δ by $N_{\text{exp}} \cdot \delta_{\text{exp}}$, where N_{exp} is the number of expansion already performed since the last successful update of Δ , before applying $N_{\text{rattle}}^{(\text{exp})}$ Monte-Carlo sweeps and applying the compression step (algorithm 2) again. We repeat this procedure with increasing δ_{exp} until either the exclusion angle is finally increased by more than the threshold or we terminate the procedure after a maximum number of expansions $N_{\text{exp}}^{(\text{max})}$. The results of the LS protocol (algorithm 1) can be reproduced by setting the trial increment of the compression step δ_0 to a small value $\delta_0 = 10^{-5}$, and by introducing some extra rattling sweeps N_{rattle} between successful increments. These extra rattling sweeps then set the timescale of compression, as N_{rattle} does in the LS protocol. The protocol is given as pseudocode in algorithm 3. Here and in the main text, whenever we refer to states produced by the softcore compression protocol, we use $N_{\text{rattle}}^{(\text{exp})} = 10$, $N_{\text{exp}} = 100$, $\delta_{\text{exp}} = 5 \times 10^{-4}$, $\delta_{\text{min}} = 10^{-6}$, $\delta_{\text{max}} = 0.05$ and $N_{\text{rattle}} = 0$.

B. Kink-free relaxation of hardcore spins

In this appendix, we give details on and present the results of softcore compression on the triangular lattice in the absence of domain wall kinks. A domain wall kink is a configuration as on the left of Fig. 3.3 (b), that is two domain walls meet at an angle of $\pi/3$ or, equivalently, three adjacent triangles have the same vector chirality. We begin by describing the adjusted preparation of initial states and subsequently discuss the results of the softcore compression protocol (algorithm 3) on the kink free model. In contrast to softcore compression in the presence of kinks, we here find [Fig. B.1 (b)] a broad distribution between $2\pi/5$ and $\pi/2$ of the exclusion angle Δ_J at which the protocol terminates for different runs, superimposed with pronounced peaks at rational fractions of 2π . The weight of the distribution shifts towards larger Δ as $N_{\text{rattle}}^{(\text{exp})}$ is increased, without ever exceeding $\Delta = \pi/2$. We interpret this as evidence for the fact that we expect the first infinite relaxation timescale of the system to occur at $\Delta = \pi/2$, where domain walls become strictly incompatible with the hardcore constraint. We attribute the presence of a broad distribution of Δ_J for small $N_{\text{rattle}}^{(\text{exp})}$ to the physics of small finite clusters and a general slowing down of dynamics due to the large Δ in conjunction with the large connectivity of the triangular lattice.

B.1. Generation of Kink-Free Initial States

We generate kink-free initial states by starting with a random state and subsequently removing kinks by a combination of a special “kink-removing” Monte-Carlo update and regular single-spin-randomization dynamics. The “kink-removing” Monte-Carlo update is shown in [Fig. B.1 (a)]. It consists of a swap of the two spins on the shorter of the two parallel edges of the parallelogram defined by the three adjacent triangles of equal chirality. Performing multiple sweeps of such updates on a random state at $\Delta = 0$ removes most, but not all domain wall kinks. We therefore alternate such sweeps with regular single-spin-randomization Monte-Carlo sweeps, but with the additional constraint that no single update should increase the total number of kinks in the system. This procedure after a few iterations successfully removes all domain wall kinks from the state, without inducing any obvious form of chiral or orientational order.

B.2. Compression of Kink-Free States

Using the kink-free initial states at $\Delta = 0$, we now run the softcore compression protocol (algorithm 3) with the additional constraint on the dynamics that no update is allowed to introduce a domain wall kink into the system.

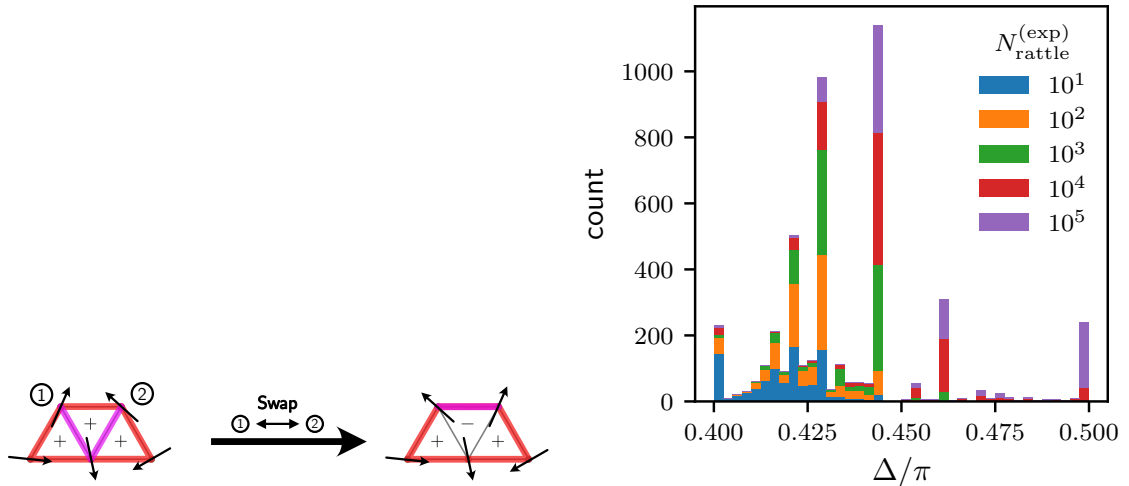


Figure B.1.: (a) Kink removal Monte-Carlo update used to generate initial states for the softcore compression in the domain wall kink free model. The sign of chirality on each triangle is indicated by $+/-$. Edges where the hardcore constraint is marginally fulfilled at $\Delta = 2\pi/5$ are indicated in red; domain walls of chirality are indicated in pink. (b) Histogram of exclusion angles Δ at which the softcore protocol (algorithm 3) for the kink free model jams, for different $N_{\text{rattle}}^{(\text{exp})}$. For $N_{\text{rattle}}^{(\text{exp})} = 10^5$, 189 of the 1000 runs reach the 120-degree state at $\Delta = 2\pi/3$.

In Fig. B.1 (b), we show a histogram of the exclusion angles Δ_J at which the protocol terminates for 10^3 initial states. First, in contrast to the model with kinks (see Fig. 3.4) here we find a broad distribution of Δ_J . We attribute this somewhat surprising result to the fact that for $\Delta > 2\pi/5$, the system finds a multitude of ways to exhibit slow dynamics with contributions from both, the nonuniversal physics of small clusters as well as the general slowing down at these exclusion angles. Note that both of these phenomena are more important on the triangular lattice than on the square lattice because of the larger connectivity of the former. The pronounced peaks on top of the broad distributions are located at rational fractions of 2π and are a consequence of what we call *Ouroboros* configurations. These are configurations on a closed loop \mathcal{C} of length L on the lattice, such that

$$\phi_j = \frac{2\pi nj}{L}, \quad j \in \mathcal{C} \quad (\text{B.1})$$

for some integer n . They are local minima of the energy functional in Eq. (4) of the main text and become forbidden at exclusion angles $\Delta = 2\pi n/L$. At these values of Δ , they introduce long but finite relaxation timescales into the system (see also discussion in the main text). These attributions are evidenced by the fact that as $N_{\text{rattle}}^{(\text{exp})}$ is increased, that is compression is performed more slowly, the broad, continuous distribution gives more and more away to a collection of peaks. Second, the weight of the distribution shifts to larger Δ_J as $N_{\text{rattle}}^{(\text{exp})}$ is increased and for $N_{\text{rattle}}^{(\text{exp})} = 10^5$, 189 of the 1000 runs even reach the 120-degree state at $\Delta = 2\pi/3$. However, no run ever fails between $\pi/2$ and $2\pi/3$. This implies that smooth domain walls of chirality are indeed the last defects to become forbidden by the hardcore constraint.

C. High-Temperature Series Expansion of the Hyperbolic Ising model

Our primary means to map out the phase diagram of the random-bond Ising model in hyperbolic space in chapter 6 is to perform high-temperature series expansions of both the susceptibility

$$\chi = \beta \frac{1}{N} \sum_{i,j} [\langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle], \quad (\text{C.1})$$

as well as of the Edwards-Anderson (EA) susceptibility

$$\chi_{\text{EA}} = \beta \frac{1}{N^2} \sum_{i,j} [\langle \sigma_i \sigma_j \rangle^2 - \langle \sigma_i \rangle^2 \langle \sigma_j \rangle^2]. \quad (\text{C.2})$$

Coming from a high-temperature, if there is a transition to low-temperatures ferromagnetic phase, the susceptibility χ at the transition should diverge as a power law

$$\chi \sim \frac{1}{(T - T_c)^\gamma} \quad (\text{C.3})$$

while the Edwards-Anderson susceptibility χ_{EA} can have either a weak singularity or diverge as well [42]. In contrast, if there is a transition into a low-temperature spin-glass phase, the susceptibility χ will exhibit only a weak singularity (a cusp), while the Edwards-Anderson susceptibility diverges as a power law

$$\chi_{\text{EA}} \sim \frac{1}{(T - T_c)^{\gamma'}}. \quad (\text{C.4})$$

C.1. Biconnected Graph Expansion of Inverse Susceptibilities

It turns out that for susceptibilities of the form

$$\chi_{k,l} = \beta \frac{1}{N} \sum_{i,j} [\langle \sigma_i \sigma_j \rangle^k - \langle \sigma_i \rangle^k \langle \sigma_j \rangle^k]^l, \quad (\text{C.5})$$

it is favourable to perform the high-temperature expansion in the *inverse* susceptibility. The reason for this is that it can be shown [272] that the only non-trivial contributions come from *biconnected* graphs, that is graphs which stay connected if any of their vertices (and the edges attached to it) are being removed. We show the




Figure C.1.: Some small biconnected subgraphs of the $\{5, 5\}$ -tiling. Removing of any vertex and all its incident edges will leave the graphs connected. Only biconnected graphs contribute to the series expansion.

first few graphs that contribute to the susceptibility $\chi = \chi_{1,1}$ and EA-susceptibility $\chi_{\text{EA}} = \chi_{2,1}$ on the $\{5, 5\}$ lattice in Fig. C.1.

The inverse susceptibility can be expanded in terms of these graphs as a function of both inverse temperature $v = \tanh(\beta J)$ and disorder strength $\mu = 1 - 2p$. In practice, the variables in the systematic biconnected graph expansion are $w = v^2$ and $\alpha = \mu/v$:

$$\tilde{\chi}^{-1}(w, \alpha) = 1 + \sum_g c(g) W(g) \quad (\text{C.6})$$

where the sum is over all graphs, $c(g)$ is the coefficient of N of the number of embeddings of the graph g into the lattice and $W(g)$ for each graph is a function of both w and α . Expanding W as a function of inverse temperature w , one can show that for each order n , the coefficient of w^n is a polynomial in α of order n with integer coefficients. For example, the inverse susceptibility on the $\{5, 5\}$ lattice is given by

$$\begin{aligned} \chi^{-1}(w, \alpha) = & 1 - 5\alpha w + 5\alpha^2 w^2 - 5\alpha^3 w^3 + 5\alpha^4 w^4 \\ & + (10\alpha + 10\alpha^2 + 10\alpha^3 + 10\alpha^4 + 5\alpha^5)w^5 \\ & + \mathcal{O}(w^6). \end{aligned} \quad (\text{C.7})$$

Note that for $\alpha = 1$ (that is $v = \mu$), we obtain the series on the Nishimori line up to order $\mathcal{O}(w^n) = \mathcal{O}(v^{2n})$.

For more details and a derivation of Eq. (C.6) see Ref. [272].

C.2. Analysis of the Series

We analyze the generates series $\tilde{\chi}(w, \alpha)$, usually for fixed α as a function of w , using *first-order homogeneous integrated differential approximants (FO-IDAs)*. One reason to choose FO-IDAs over simpler methods is that they are known to be less biased towards the lower-order coefficients of the expansion [273]. This is important, as the most relevant contributions on a $\{r, s\}$ tiling come from graphs with at least r edges.

The analysis using FO-IDAs proceeds as follows: For fixed disorder strength α , we assume that the series $\tilde{\chi}$ is the solution of a first-order differential equation of

the form

$$Q_L(w) \frac{d\tilde{\chi}(w)}{dv} + R_M(w) \tilde{\chi}(w) + S_T(w) = 0 \quad (\text{C.8})$$

where Q_L , R_M and S_T are polynomials of degree L , M , T , respectively. By equating the series order-by-order with the coefficients of Eq. (C.8) we obtain a linear system of equations in the coefficients of the polynomials Q_L , R_M and S_T . It can be shown that for any root w_c of the polynomial Q_L , a solution of Eq. (C.8) has an algebraic singularity of the form $(w - w_c)^{-\gamma}$ [184]. The exponent of the singularity is given by

$$\gamma = \frac{R_M(w_c)}{Q'_L(w_c)}. \quad (\text{C.9})$$

Generally, the results for w_c and γ will depend on the choice of degrees L , M and T . If we have the series up to order N then we can choose all possible values satisfying $L + M + T \leq N - 2$. Following [273] we exclude approximants if one of the following is true

- a root of R_M is close to w_c , giving rise to a small estimate of γ
- a complex root of Q_L with small absolute value smaller than w_c is close to the real axis

We observe that the convergence of the series is very good, since the approximants for different choices of L , M and T are all close.

D. Neutron Scattering Cross Section of Dipolar-Octupolar Pyrochlores

Here, we derive the full expression for the neutron scattering cross section of dipolar octupolar pyrochlore materials.

$$\left(\frac{d\sigma}{d\Omega}\right)_{\text{mag}} = (\gamma r_0)^2 S(\mathbf{q}) \quad (\text{D.1})$$

where $\gamma r_0 \approx 5.39 \times 10^{-15}$ m is the magnetic scattering length and $S(\mathbf{q})$ is the magnetic structure factor. It is defined as

$$S(\mathbf{q}) = \sum_{\alpha, \beta} \left(\delta_{\alpha\beta} - \frac{q^\alpha q^\beta}{q^2} \right) \langle \overline{M^\alpha}(\mathbf{q}) \cdot M^\beta(\mathbf{q}) \rangle \quad (\text{D.2})$$

where $\overline{}$ denotes complex conjugation and $\mathbf{M}(\mathbf{q})$ denotes the Fourier transform of the magnetization operator.

The main result of this appendix will be to express the structure factor in Eq. (D.2) in terms of the pseudospin operators S_j^α appearing in the exchange Hamiltonian [Eq. (7.27)].

D.1. Matrix Elements of Magnetization Operator

To this end, consider the matrix elements of the magnetization operator \mathbf{M} with respect to the electron states that the pseudospin operators act on. Treating the wave function of the 4f electronic states as hydrogen-like with effective nuclear charge Z_{eff} yields [219, 249–251]

$$\langle SLJM | M^\alpha(\mathbf{k}) | SLJM' \rangle = \sqrt{4\pi} \sum_{p=1}^{2J} \frac{Z_p(k)}{p} \Phi_p^\alpha(\hat{\mathbf{k}}). \quad (\text{D.3})$$

The radial part is given in terms of expectation values of spherical Bessel functions

$$Z_p(k) = c_{p-1} \langle j_{p-1}(k) \rangle + c_{p+1} \langle j_{p+1}(k) \rangle \quad (\text{D.4a})$$

and the coefficients c_p depend on the material considered. They have been tabulated for different rare-earth ions for example in Ref. [274] (see pp. 160f for Ce^{3+}) and the values for the $J = \frac{5}{2}$ doublet of Ce^{3+} [Eq. (7.21)] are reproduced in Table D.1. The radial expectation value of the spherical Bessel function j_p for a hydrogen-like electron density with principal quantum number n and angular momentum quantum number L ($n = 4$ and $L = 3$ for 4f electrons) are given explicitly by

$$\langle j_p(k) \rangle = \int_0^\infty j_p(kr) r^2 R_{n,L}(r) dr \quad (\text{D.4b})$$

p	0	2	4	6
c_p	$\frac{15}{14}$	$\frac{3}{14\sqrt{5}}$	$\frac{5}{1386}$	0

Table D.1.: Coefficients c_p in Eq. (D.4) for Ce^{3+} Ions assuming a pure $J = \frac{5}{2}$ doublet. For values for other ions see also pp. 160f of Ref. [274].

with the radial part of the hydrogen-like wave function

$$R_{n,L}(r) = \sqrt{\left(\frac{2Z_{\text{eff}}}{na_0}\right)^3 \frac{(n-l-1)!}{2n(n+l)!}} \exp\left(-\frac{2Z_{\text{eff}}}{na_0}r\right) \left(\frac{2Z_{\text{eff}}}{na_0}\right)^L L_{n-L-1}^{(2L+1)}\left(\frac{2Z_{\text{eff}}}{na_0}r\right), \quad (\text{D.4c})$$

where $a_0 \approx 0.529177211 \text{ \AA}$ is the Bohr radius and $L_{n-L-1}^{(2L+1)}$ are the generalized Laguerre polynomials.

The angular dependence is given in terms of spherical harmonics as

$$\Phi_p^x(\hat{\mathbf{k}}) = \sum_{q=-p}^p \Pi(p, q, J, M, M') \frac{1}{2} \left[Y_{p-1}^{q+1}(\hat{\mathbf{k}}) \sqrt{(p-q)(p-q-1)} - Y_{p-1}^{q-1}(\hat{\mathbf{k}}) \sqrt{(p+q)(p+q-1)} \right] \quad (\text{D.5a})$$

$$\Phi_p^y(\hat{\mathbf{k}}) = \sum_{q=-p}^p \Pi(p, q, J, M, M') \frac{1}{2} \left[Y_{p-1}^{q+1}(\hat{\mathbf{k}}) \sqrt{(p-q)(p-q-1)} + Y_{p-1}^{q-1}(\hat{\mathbf{k}}) \sqrt{(p+q)(p+q-1)} \right] \quad (\text{D.5b})$$

$$\Phi_p^z(\hat{\mathbf{k}}) = \sum_{q=-p}^p \Pi(p, q, J, M, M') Y_{p-1}^q(\hat{\mathbf{k}}) \sqrt{(p-q)(p+q)} \quad (\text{D.5c})$$

where $Y_l^m(\theta, \phi)$ are spherical harmonics. The argument of these spherical harmonics importantly is evaluated in the *local* basis, that is the angles θ and ϕ for each of the four sublattices are taken with respect to a local coordinate system $\hat{\mathbf{x}}_j, \hat{\mathbf{y}}_j, \hat{\mathbf{z}}_j$. The choice of this local coordinate system here is *not* arbitrary and in the above expression we use

$$\hat{\mathbf{x}}_1 = \begin{pmatrix} -1 \\ -1 \\ 2 \end{pmatrix} \quad \hat{\mathbf{y}}_1 = \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} \quad \hat{\mathbf{z}}_1 = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \quad (\text{D.5d})$$

$$\hat{\mathbf{x}}_2 = \begin{pmatrix} -1 \\ 1 \\ -2 \end{pmatrix} \quad \hat{\mathbf{y}}_2 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} \quad \hat{\mathbf{z}}_2 = \begin{pmatrix} 1 \\ -1 \\ -1 \end{pmatrix}, \quad (\text{D.5e})$$

$$\hat{\mathbf{x}}_3 = \begin{pmatrix} 1 \\ -1 \\ -2 \end{pmatrix} \quad \hat{\mathbf{y}}_3 = \begin{pmatrix} -1 \\ -1 \\ 0 \end{pmatrix} \quad \hat{\mathbf{z}}_3 = \begin{pmatrix} -1 \\ 1 \\ -1 \end{pmatrix}, \quad (\text{D.5f})$$

$$\hat{\mathbf{x}}_4 = \begin{pmatrix} 1 \\ 1 \\ 2 \end{pmatrix} \quad \hat{\mathbf{y}}_4 = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix} \quad \hat{\mathbf{z}}_4 = \begin{pmatrix} -1 \\ -1 \\ 1 \end{pmatrix}. \quad (\text{D.5g})$$

The factor Π is defined in terms of $3j$ symbols

$$\Pi(p, q, J, M, M') = (-1)^{J-M'} \frac{\begin{pmatrix} p & J & J \\ q & M & -M' \end{pmatrix}}{\begin{pmatrix} p & J & J \\ 0 & J & -J \end{pmatrix}} \quad (\text{D.5h})$$

which enforce $q = M' - M$.

D.2. Structure Factor

Using the expression for the matrix elements of the magnetization operator in Eq. (D.3) and below, we can write this operator as a two-by-two matrix in the subspace spanned by the crystal field ground state doublet. Since \mathbf{M} is traceless it is then possible to write it as a combination of pseudospin operators in this space

$$M_j^\alpha = \sum_{\gamma} \mathcal{D}_j^{\alpha\gamma} S_j^\gamma \quad (\text{D.6})$$

and we call the prefactors $D_j^{\alpha\gamma}(\mathbf{q})$ the generalized form factor. Note that because of the dependence of the matrix elements on the local basis, the above expression still carries a sublattice index j . As already explained in the main text one can then express the structure factor in terms of the pseudospin operators as

$$S(\mathbf{q}) = \sum_{ij} e^{-i\mathbf{q}(\mathbf{r}_i - \mathbf{r}_j)} \sum_{\alpha\beta\gamma\delta} \left(\delta_{\alpha\beta} - \frac{q^\alpha q^\beta}{q^2} \right) \mathcal{D}_i^{\alpha\gamma}(\mathbf{q}) \mathcal{D}_j^{\beta\delta}(\mathbf{q}) C_{ij}^{\gamma\delta}(\mathbf{q}) \quad (\text{D.7})$$

where i and j enumerate the four pyrochlore sublattices at positions \mathbf{r}_i and \mathbf{r}_j respectively and

$$C_{ij}^{\gamma\delta}(\mathbf{q}) = \sum_{\mathbf{x}, \mathbf{x}'} e^{-i\mathbf{q}(\mathbf{x} - \mathbf{x}')} \langle S_i^\alpha(\mathbf{x}) S_i^\beta(\mathbf{x}') \rangle \quad (\text{D.8})$$

is the Fourier transform of the two-point correlation function between pseudospin components γ, δ on sublattices i and j . The sum above is over lattice vectors \mathbf{x} .

D.3. Explicit Expression for Ce^{3+} Doublet

The generalized form factor in general is a sum over multipole orders p , which contribute at increasingly large momentum since for small k , $Z_p(k) \sim k^{p-1}$. It is convenient to write the matrix \mathcal{D}_j as a sum over individual components

$$\mathcal{D}_j = \sum_{p=1}^{2J} \sum_{\alpha=x,y,z} \hat{\alpha}_j \otimes \mathbf{T}_{p,j}^{(\alpha)} \quad (\text{D.9})$$

where $\hat{\mathbf{x}}_j, \hat{\mathbf{y}}_j, \hat{\mathbf{z}}_j$ are the local basis vectors for each sublattice given in Eq. (D.5) and \otimes denotes the outer product between vectors. In this case

$$M_j^\alpha = \sum_p \sum_{\alpha=x,y,z} \hat{\alpha}_j \left(\mathbf{T}_{p,j}^{(\alpha)} \cdot \mathbf{S}_j \right) \quad (\text{D.10})$$

where $\hat{\mathbf{x}}_j, \hat{\mathbf{y}}_j, \hat{\mathbf{z}}_j$ are the local basis vectors for each sublattice given in Eq. (D.5). In the following, we treat Ce as a hydrogen atom with effective nuclear charge $Z_{\text{eff}} = 17$ and assume a pure $J = \frac{5}{2}$ doublet.

Dipole contribution

The first order $p = 1$ is the well-known dipolar contribution [216, 217]

$$\mathbf{T}_{p,j}^{(x)}(\boldsymbol{\kappa}) = \mathbf{T}_{p,j}^{(y)}(\mathbf{q}) = 0, \quad (\text{D.11})$$

$$\mathbf{T}_{p,j}^{(z)}(\boldsymbol{\kappa}) = \frac{g_z}{2} f_1(\kappa) \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad (\text{D.12})$$

where $\boldsymbol{\kappa} = 2\mathbf{q}a_0/(a Z_{\text{eff}})$ with the Bohr radius a_0 and lattice constant a (2.6629 Å for Ce). The proportionality factor $g_z = 18/7 \approx 2.57$ is the g-factor of the Ce ground state doublet. The radial integral for a hydrogen-like wavefunction can be solved explicitly which yields

$$f_1(\kappa) = \frac{35 + 364\kappa^2 - 3760\kappa^4 + 2880\kappa^6}{35(1 + 4\kappa^2)^8}, \quad (\text{D.13a})$$

where $\kappa = |\boldsymbol{\kappa}|$. Note that the dipolar order of the full scattering amplitude as calculated here it not in exact agreement with the dipolar approximation to the form factor [251]. For all practical purposes, the difference between Eq. (D.12) and the dipolar approximation are however negligible. This verifies the applicability of the latter in situations where the local moments are dipoles.

Octupole contribution

The next nonzero order in the multipole expansion corresponds to the octupole moment $p = 3$ and is given by

$$\mathbf{T}_{3,j}^{(x)}(\boldsymbol{\kappa}) = f_3(\kappa) \begin{pmatrix} \frac{36}{7} [\kappa_x^2 - \kappa_y^2] \\ -\frac{72}{7} \kappa_x \kappa_y \\ -\frac{36}{5} \kappa_x \kappa_z \end{pmatrix} \quad (\text{D.14})$$

$$\mathbf{T}_{3,j}^{(y)}(\boldsymbol{\kappa}) = f_3(\kappa) \begin{pmatrix} -\frac{72}{7} \kappa_x \kappa_y \\ -\frac{36}{7} [\kappa_x^2 - \kappa_y^2] \\ -\frac{36}{5} \kappa_y \kappa_z \end{pmatrix} \quad (\text{D.15})$$

$$\mathbf{T}_{3,j}^{(z)}(\boldsymbol{\kappa}) = f_3(\kappa) \begin{pmatrix} 0 \\ 0 \\ \frac{18}{5} [\kappa^2 - 3\kappa_z^2] \end{pmatrix} \quad (\text{D.16})$$

where

$$f_3(\kappa) = \frac{9 + 200\kappa^2 - 240\kappa^4}{9(1 + 4\kappa^2)^8} \quad (\text{D.17})$$

and $\kappa_\alpha = \boldsymbol{\kappa} \cdot \hat{\boldsymbol{\alpha}}$ for $\alpha = x, y, z$ and $\hat{\boldsymbol{x}}_j, \hat{\boldsymbol{y}}_j, \hat{\boldsymbol{z}}_j$ are the local basis vectors for each sublattice given in Eq. (D.5).

Triakontadipole contribution

The final nonzero order when assuming a pure $J = \frac{5}{2}$ ground state is the triakontadipole ($p = 5$). It is given by

$$\mathbf{T}_{5,j}^{(x)}(\boldsymbol{\kappa}) = f_5(\kappa) \begin{pmatrix} \frac{40}{21} (\kappa_x^4 + \kappa_y^4 - 6\kappa_x^2\kappa_y^2 - 4[\kappa_x^2 - \kappa_y^2][7\kappa_z^2 - \kappa^2]) \\ \frac{160}{21}\kappa_z\kappa_y [14\kappa_z^2 - 2\kappa^2 - \kappa_x^2 + \kappa_y^2] \\ -\frac{800}{147}\kappa_z\kappa_x [7\kappa_z^2 - 3\kappa^2] \end{pmatrix} \quad (D.18)$$

$$\mathbf{T}_{5,j}^{(y)}(\boldsymbol{\kappa}) = f_5(\kappa) \begin{pmatrix} \frac{160}{21}\kappa_z\kappa_y [14\kappa_z^2 - 2\kappa^2 + \kappa_x^2 - \kappa_y^2] \\ \frac{40}{21} (\kappa_x^4 + \kappa_y^4 - 6\kappa_x^2\kappa_y^2 + 4[\kappa_x^2 - \kappa_y^2][7\kappa_z^2 - \kappa^2]) \\ -\frac{800}{147}\kappa_z\kappa_y [7\kappa_z^2 - 3\kappa^2] \end{pmatrix} \quad (D.19)$$

$$\mathbf{T}_{5,j}^{(z)}(\boldsymbol{\kappa}) = f_5(\kappa) \begin{pmatrix} \frac{640}{21}\kappa_z [\kappa_x^3 - 3\kappa_x\kappa_y^2] \\ \frac{640}{21}\kappa_z [\kappa_y^3 - 3\kappa_y\kappa_x^2] \\ -\frac{25}{147} [40\kappa_z^2\kappa^2 - 11\kappa^4 + 35(\kappa_z^4 - 6\kappa_z^2[\kappa_x^2 + \kappa_y^2]) + (\kappa_x^2 + \kappa_y^2)^2] \end{pmatrix} \quad (D.20)$$

where

$$f_5(\kappa) = \frac{1 + 60\kappa^2}{(1 + 4\kappa^2)^8} \quad (D.21)$$

and κ_α defined as before.

E. Semiclassical Monte-Carlo Molecular Dynamics

This appendix provides a brief but self-contained introduction to semiclassical Monte-Carlo molecular dynamics simulations and how to use them to obtain dynamical properties of spin systems.

E.1. Landau-Lifshitz Dynamics

The general idea of semiclassical Monte-Carlo molecular dynamics (MCMD) simulations is quickly explained. One first generates an ensemble of classical equilibrium states of the system at hand, e.g. by using classical Monte Carlo simulations and then evolves these states in time using the semiclassical Landau-Lifshitz equations [275]

$$\frac{d}{dt}\mathbf{S}_i = \mathbf{S}_i \times \mathbf{H}_i \quad (\text{E.1})$$

where $\mathbf{H}_i = d\mathcal{H}/d\mathbf{S}_i$ is the local field at site i and \mathcal{H} is the Hamiltonian of the system. Dynamical properties are then calculated as averages over the ensemble of initial states.

In the following, we will explain the method in detail using as an illustrative example the pyrochlore Heisenberg antiferromagnet (PHAF)

$$\mathcal{H} = J \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{S}_j, \quad (\text{E.2})$$

where the sum is over nearest-neighbor bonds $\langle ij \rangle$ in the pyrochlore lattice. The dynamics of this model has first been studied using MCMD by Conlon and Chalker [276]. We note that this method has also been applied to the Heisenberg model on the Kagome lattice [277] and to a slightly perturbed version on the pyrochlore lattice [254]. In Ref. [254] it was even shown to reproduce the experimentally measured dynamical structure factor of $\text{NaCaNi}_2\text{F}_7$ with remarkable accuracy. It was also used in Ref. [226] to obtain the structure factor of a model fitted to describe $\text{Ce}_2\text{Zr}_2\text{O}_2$, the same material discussed in the main text in chapter 8 and chapter 9.

The method as presented here generalizes readily to the anisotropic model XYZ model studied in the main text.

E.2. Generating the Initial States

The initial states for the MCMD simulation can be generated by any method that produces a thermal ensemble. An easy and very versatile way to do exactly that are classical Monte Carlo simulations. For the Heisenberg model at hand and also

for the anisotropic model studied in the main text, we use the exact heat bath algorithm [278]. The autocorrelation time of exact heat bath MC has the same scaling ($\tau \sim T^{-1}$) as other local algorithms such as metropolis, but it still provides a significant parametric speedup at low temperatures. We then use snapshots of these independent Monte Carlo simulations as initial states for the molecular dynamics simulations.

E.3. Molecular Dynamics Evolution

As discussed before, the semiclassical time evolution of a spin state $\{\mathcal{S}_i\}$ is given by the Landau-Lifshitz [275] equation, which describes precession of each spin around their local magnetic field

$$\frac{d}{dt}\mathcal{S}_i = \mathcal{S}_i \times J \sum_{j \in \text{nn}(i)} \mathcal{S}_j. \quad (\text{E.3})$$

In practice, this nonlinear equation is straightforwardly integrated using a fourth-order Runge-Kutta method. We note that Eq. (E.3) conserves both the total energy and magnetization of the system, which can be checked to ensure that the numerical error is controlled. For the PHAF, choosing a time step $\Delta t = 10^{-4}J^{-1}$ preserves both energy and all three components of the magnetization up to machine precision even up to times as long as $1000J^{-1}$.

We are most interested in computing the dynamic structure factor, which is the power spectral density of the dynamic correlation function

$$S(\omega, \mathbf{q}) = \int_{-\infty}^{\infty} dt \sum_{m,n} e^{i\omega t} e^{-i\mathbf{q} \cdot (\mathbf{r}_i - \mathbf{r}_j)} \langle \mathcal{S}_i(t) \cdot \mathcal{S}_j \rangle, \quad (\text{E.4})$$

$$= |\mathcal{F}[\mathcal{S}_i(t)]|^2, \quad (\text{E.5})$$

where \mathcal{F} denotes the Fourier transform in all spacial direction as well as in the time direction. The last equality is also known as the Wiener-Khinchin theorem and follows from the convolution theorem [279].

In practice, we compute $S(\omega, \mathbf{q})$ from the time-evolved spin configuration $\{\mathcal{S}_i(t)\}$ by computing the Fourier transform in all three space directions using periodic boundary conditions. In the time direction, since we cannot enforce periodic boundary conditions, we take a periodogram with a Hann window. This is important to avoid finite-time effects as explained in the following.

One method of estimating the PSD using a finite time-slice is by taking a periodogram, that is taking the Fourier transform magnitude squared of the spin configuration $\mathcal{S}_i(t)$ times a window function $W(t)$

$$|\mathcal{F}[W(t)\mathcal{S}_i(t)]|^2 = |\mathcal{F}[W(t)] \star \mathcal{F}[\mathcal{S}_i(t)]|^2 \quad (\text{E.6})$$

where \star denotes the convolution. The 'trivial' solution of taking the FFT magnitude squared of the spin configuration on the finite time slice directly is equivalent to taking a periodogram using the 'boxcar' window

$$W_{\text{BS}}(t) = \begin{cases} 1 & t \in [0, T_{\text{fin}}] \\ 0 & \text{else} \end{cases}. \quad (\text{E.7})$$




Figure E.1.: Estimation of the structure factor from a finite time slice of length $T_{\text{fin}} = 200J^{-1}$ (a, c) and $T_{\text{fin}} = 1000J^{-1}$ (b, d), by taking a periodogram without a window (a,b) and using the Hann window (c,d). Along the path $[l, l, l]$, the elastic structure factor $S(0, \mathbf{q})$ is nonzero only at $l = 1/2$.

Regardless of the window used, the result of this procedure is the convolution of the dynamic structure with the Fourier transform of the window $W(t)$.

For $T_{\text{fin}} \rightarrow \infty$, the Fourier transform of the boxcar window becomes a delta peak at $\omega = 0$, and we recover the Wiener-Khinchin theorem. If the final time T_{fin} is not large enough however, the slow decay of $\mathcal{F}[W_{\text{BC}}(t)]$ can lead to significant finite time effects especially at finite energy, since typically the zero energy component of the structure factor is much larger than its part at finite-energy. Broadening this zero-energy component significantly will make it 'leak' into the finite energy part of $S(\omega, \mathbf{q})$.

This is shown for example in Fig. E.1 for the Heisenberg model on the pyrochlore lattice. We simulate a system with $N = 16384$ spins at $T = 0.010J$ and estimate the dynamical structure factor from a finite time slice. Without proper windowing for a time slice as long as $T_{\text{fin}} = 200J^{-1}$ spurious features can be introduced into the signal as shown in panel (a).

E.4. Spectral Classical-to-Quantum Conversion

The molecular dynamics solution as described above computes the *classical* dynamics of the model. It treats the spins as classical magnetic moments precessing in

their local field. To compare this to the (quantum) experiment or a theoretical method such as linear spin wave theory, one has to re-scale the classical calculation by some factor. This is because the classical dynamical structure factor is symmetric with respect to neutron energy-transfer E , and it vanishes as T approaches zero for all $E > 0$. Neither of these is the case for the dynamical structure factor of the quantum system. Another, more quantitative, way to think about this is via the fluctuation-dissipation theorem by comparing the version for classical and quantum systems [280]. In particular, for a classical system we get

$$(\beta E)S_{\text{classical}}(\mathbf{Q}, E, T) = \text{Im}[\chi''(\mathbf{Q}, E, T)] \quad (\text{E.8})$$

while for the quantum system it reads

$$(1 - e^{-\beta E})S_{\text{quantum}}(\mathbf{Q}, E, T) = \text{Im}[\chi''(\mathbf{Q}, E, T)], \quad (\text{E.9})$$

where $\beta = 1/(k_B T)$. It is then reasonable to equate the imaginary part of the susceptibility, $\text{Im}[\chi''(\mathbf{Q}, E, T)]$, as this quantity is real and symmetric for both the classical and the quantum system. Furthermore, as shown in Ref. [254], $\text{Im}[\chi''_{\text{quantum}}] = \text{Im}[\chi''_{\text{classical}}]$ within linear spin wave theory. Using the quantum and classical fluctuation dissipation theorem for the respective sides then yields

$$S_{\text{quantum}}(\mathbf{Q}, E, T) = \frac{\beta E}{1 - e^{-\beta E}} S_{\text{classical}}(\mathbf{Q}, E, T), \quad (\text{E.10})$$

which is what we use to estimate the dynamical structure factor of the (quantum) experiment using our classical simulation. The dynamical structure factor is then powder-averaged to obtain $S_{\text{quantum}}(|\mathbf{Q}|, E, T)$, and convolved with the experimental resolution.

Note that Eq. (E.10) accounts for detailed balance

$$S_{\text{quantum}}(\mathbf{Q}, -E, T) = e^{-\beta E} S_{\text{quantum}}(\mathbf{Q}, E, T) \quad (\text{E.11})$$

since $S_{\text{classical}}(\mathbf{Q}, E, T) = S_{\text{classical}}(\mathbf{Q}, -E, T)$. Zhang *et al.* (Ref. [254]) derive the conversion factor βE by comparing the classical spin wave theory at finite temperature with the quantum spin wave theory at zero temperature. It is thus valid in the case $\beta E \gg 1$, which is well fulfilled in their case, but not applicable in the general case. However, note that the factor $\frac{\beta E}{1 - e^{-\beta E}}$ reduces to βE for $\beta E \gg 1$, so the calculation presented here is entirely consistent with the argument in Ref. [254].

F. Real Space Perturbation Theory of Chain Phases in Dipolar-Octupolar Pyrochlores

In this appendix, we derive the expression for the leading-order effective chain couplings presented in Sec. 9.2.1 of chapter 9 using real space perturbation theory. We start from a Hamiltonian of the form

$$\mathcal{H} = \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{J} \cdot \mathbf{S}_j - \sum_i \mathbf{h}_i \cdot \mathbf{S}_i, \quad (\text{F.1})$$

from which Eq. (9.1) is obtained for the choice $\mathbf{J} = \text{diag}(J_x, J_y, J_z)$ and $\mathbf{h}_i = \mu_B g_z \mathbf{H} \cdot \hat{\mathbf{z}}_i (\sin \theta, 0, \cos \theta)$, with $\mathbf{H} \parallel (1, 1, 0)$.

We transform to a local basis, $\hat{\mathbf{u}}_i, \hat{\mathbf{v}}_i, \hat{\mathbf{w}}_i$ chosen such that $\mathbf{S}_i^{(0)} = S \hat{\mathbf{w}}_i$ is a classical ground state. The Hamiltonian is then separated an unperturbed part, \mathcal{H}_0 and four transverse perturbations

$$\mathcal{H} = \mathcal{H}_0 + \sum_{\langle ij \rangle} V_{ij}^{(1)} + V_{ij}^{(2)} + V_{ij}^{(3)} + V_{ij}^{(4)}, \quad (\text{F.2})$$

where

$$\mathcal{H}_0 = E_0 + \sum_i B_i \delta S_i, \quad (\text{F.3a})$$

$$E_0 = \frac{S^2}{2} \sum_i \left(\sum_{j \in \text{nn}(i)} \hat{\mathbf{w}}_i \cdot \mathbf{J} \cdot \hat{\mathbf{w}}_j \right) - S \hat{\mathbf{w}}_i \cdot \mathbf{h}_i, \quad (\text{F.3b})$$

$$B_i = -S \sum_{j \in \text{nn}(i)} \mathbf{J} \cdot \hat{\mathbf{w}}_j + \mathbf{h}_i, \quad (\text{F.3c})$$

and

$$V_{ij}^{(1)} = \mathbf{c}_i^* \cdot \mathbf{J} \cdot \mathbf{c}_j^* S_i^+ S_j^+ + h.c., \quad (\text{F.4a})$$

$$V_{ij}^{(2)} = \mathbf{c}_i^* \cdot \mathbf{J} \cdot \mathbf{c}_j S_i^+ S_j^- + h.c., \quad (\text{F.4b})$$

$$V_{ij}^{(3)} = \hat{\mathbf{w}}_i \cdot \mathbf{J} \cdot \hat{\mathbf{w}}_j \delta S_i \delta S_j, \quad (\text{F.4c})$$

$$V_{ij}^{(4)} = \mathbf{c}_i \cdot \mathbf{J} \cdot \hat{\mathbf{w}}_i S_i^- \delta S_j + \hat{\mathbf{w}}_i \cdot \mathbf{J} \cdot \mathbf{c}_j \delta S_i S_j^- + h.c., \quad (\text{F.4d})$$

where $\delta S_i = S - S_i^w$ and $\mathbf{c}_i = \frac{1}{2} (\hat{\mathbf{u}}_i + i \hat{\mathbf{v}}_i)$.

Going forward, we will first calculate the effective couplings in the Chain_Z phase (see Sec. 9.1), with the corresponding results for the Chain_X phase being obtained by the same procedure. Finally, the effective couplings in the Chain_Y phase will be computed in the last part of this appendix.

For a ground state in the Chain_Z phase (see Sec. 9.1) we have

$$\hat{\mathbf{w}}_i = \begin{cases} (\sigma_i^\alpha \sin \phi, 0, \sigma_i^\alpha \cos \phi) & i \in \alpha \text{ chain} \\ (0, 0, \sigma_i^\beta \eta_c) & i \in \beta \text{ chain} \end{cases}, \quad (\text{F.5a})$$

and choose

$$\mathbf{c}_i = \frac{1}{2} \begin{cases} (\sigma_i^\alpha \cos \phi, i, -\sigma_i^\alpha \sin \phi) & i \in \alpha \text{ chain} \\ (1, i\sigma_i^\beta \eta_c, 0) & i \in \beta \text{ chain} \end{cases}, \quad (\text{F.5b})$$

where η_c is the Ising variable of the chain that the site i belongs to

$$\sigma_i^\alpha = \begin{cases} 1 & \text{if } i \in \mathcal{L}_0 \\ -1 & \text{if } i \in \mathcal{L}_3 \end{cases}, \quad (\text{F.6a})$$

$$\sigma_i^\beta = \begin{cases} 1 & \text{if } i \in \mathcal{L}_1 \\ -1 & \text{if } i \in \mathcal{L}_2 \end{cases}. \quad (\text{F.6b})$$

The angle ϕ above parameterizes the classical configuration of spins on the α chains such that $\theta - \phi$ is the angle between the local field and the magnetic moments. ϕ thus depends on the external field h, θ as well as on the couplings J_x, J_z and $\phi \rightarrow \theta$ as $h \rightarrow \infty$.

The perturbations $V_{ij}^{(1\dots 4)}$ hence depend parametrically on the chain configuration $\{\eta_i\}$ via the basis Eq. (F.5). Standard perturbation theory yields a configuration dependent energy correction $\Delta E(\{\eta_i\})$. As we will see below, the leading order contribution to ΔE depends only on products of nearest-neighbor chains such that

$$\Delta E = \sum_{\langle ij \rangle} \delta E_{ij} = J_{\text{ch}} \sum_{\langle ij \rangle} \eta_i \eta_j + J'_{\text{ch}} \sum_{\langle ij \rangle'} \eta_i \eta_j, \quad (\text{F.7})$$

which yields the effective Ising model discussed in Sec. 9.2.1.

F.1. Perturbation Theory

Since the energy corrections must stay extensive, any term that contributes to perturbation theory must correspond to a linked cluster of the pyrochlore lattice, where each edge represents one or more applications of the perturbation $V = V^{(1)} + V^{(2)} + V^{(3)} + V^{(4)}$. The energy correction corresponding to a linked cluster with p edges is [263]

$$\delta E^{(p)} = \sum_{\{\psi_i\}} \frac{\langle 0 | V | \psi_1 \rangle \langle \psi_1 | V | \psi_2 \rangle \dots \langle \psi_{p-1} | V | 0 \rangle}{(E_0 - E_{\psi_1}) \dots (E_0 - E_{\psi_{p-1}})}, \quad (\text{F.8})$$

where E_0 is the classical ground state energy and E_{ψ_i} is the unperturbed energy of the intermediate state $|\psi_i\rangle$. From the above we can read off some simple rules:

1. Any linked cluster corresponds to a sum of terms, which are all possible combinations of applications of the $V_{ij}^{(1\dots 4)}$ to the links such that we begin and end in the vacuum of spin flips $|0\rangle$.

2. Any such sequence of perturbations has to begin and end with either $V^{(1)}$ or $V^{(4)}$.
3. Any such sequence must involve an even number of applications of $V^{(4)}$.

Since we are interested only in the leading order contribution to the energy correction we do not have to consider the perturbations $V^{(3)}$ and $V^{(4)}$ since the operator δS_i does not modify the spin configuration. Thus for any sequence of applications of perturbations involving $V^{(3)}$ or $V^{(4)}$, there will be a nonvanishing lower order term that contributes.

The matrix elements of the perturbations are evaluated straightforwardly using Eq. (F.4) and Eq. (F.5). If i, j are both on a β chain

$$V_{ij}^{(1)} = \frac{1}{4} (J_x + J_y) S_i^+ S_j^+ + h.c., \quad (\text{F.9a})$$

$$V_{ij}^{(2)} = \frac{1}{4} (J_x - J_y) S_i^+ S_j^- + h.c.. \quad (\text{F.9b})$$

In contrast, if i, j are both on an α chain

$$V_{ij}^{(1)} = -\frac{1}{4} (\cos \phi^2 J_x + J_y + \sin \phi^2 J_z) S_i^+ S_j^+ + h.c., \quad (\text{F.9c})$$

$$V_{ij}^{(2)} = \frac{1}{4} (-\cos \phi^2 J_x + J_y - \sin \phi^2 J_z) S_i^+ S_j^- + h.c.. \quad (\text{F.9d})$$

Finally, if i is on a β chain and j is on an α chain

$$V_{ij}^{(1)} = \frac{1}{4} (\sigma_j^\alpha \cos \phi J_x - \sigma_i^\beta \eta_c J_y) S_i^+ S_j^+ + h.c., \quad (\text{F.9e})$$

$$V_{ij}^{(2)} = \frac{1}{4} (\sigma_j^\alpha \cos \phi J_x + \sigma_i^\beta \eta_c J_y) S_i^+ S_j^- + h.c.. \quad (\text{F.9f})$$

Similarly, the factors in the denominator of Eq. (F.8) are given by

$$\begin{aligned} E_0 - E_\psi &= \langle \psi | E_0 - H_0 | \psi \rangle \\ &= - \sum_i B_i \langle \psi | \delta S_i | \psi \rangle, \end{aligned} \quad (\text{F.10})$$

where if i is on a β chain:

$$B_i = 2S J_z, \quad (\text{F.11a})$$

and if i is on an α chain:

$$B_i = 2S (\sin(\phi)^2 J_x + \cos(\phi)^2 J_z) + \cos(\phi - \theta) h. \quad (\text{F.11b})$$

F.2. Degeneracy Lifting

We are interested in the lowest order quantum correction that lifts the classical degeneracy. For that, the energy correction must be proportional to the product of

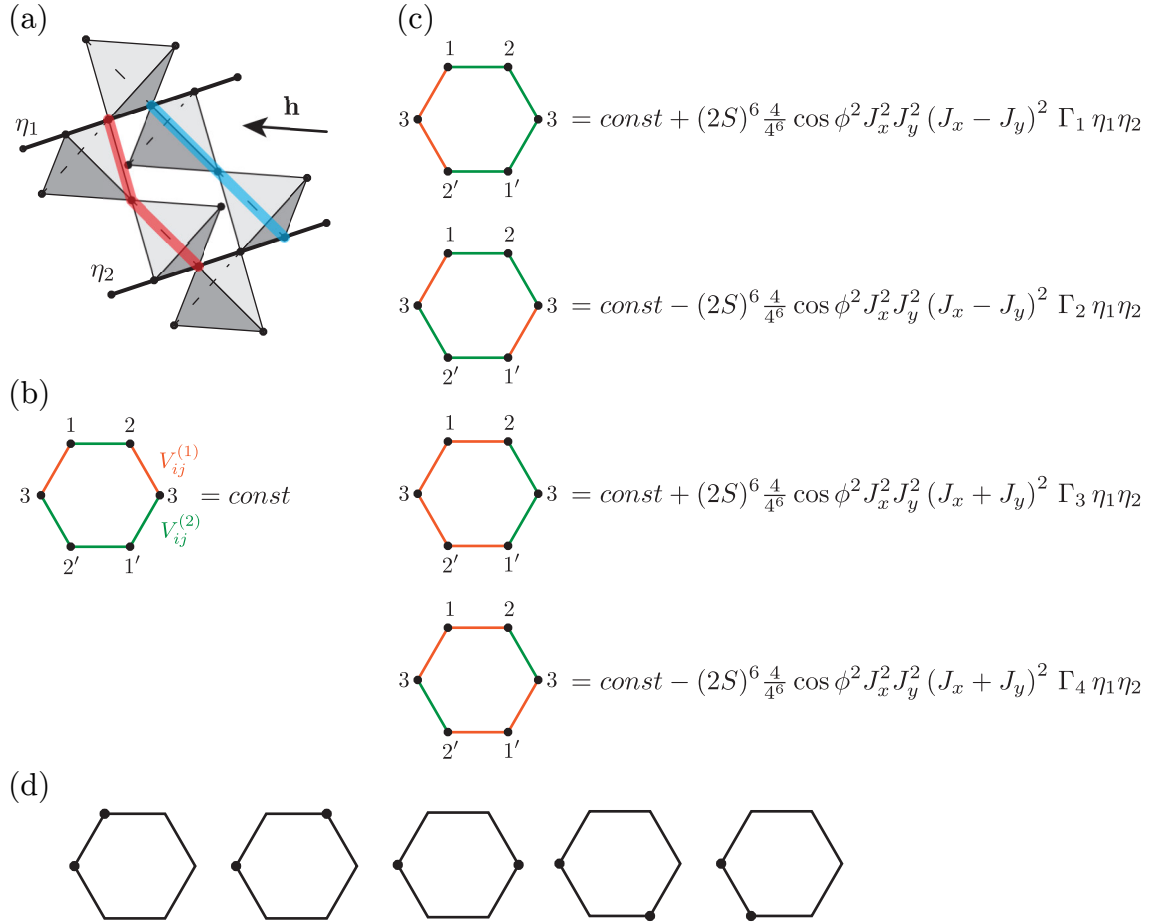


Figure F.1.: (a) The two distinct smallest clusters connecting two chains, marked in red and blue. Their lowest order contribution is of order 4 and they cancel out exactly. (b) A configuration of perturbation on a hexagon that does not contribute to the effective coupling. (c) The four inequivalent configurations that do contribute to J_{ch} . (d) A possible sequence of intermediate states for the first configuration in (c). In (b) and (c), we label the sites with reference to the sublattice \mathcal{L}_ν they belong to, consistent with the hexagon indicated in green in Fig. 9.5 (a)

as least two chains $\eta_i\eta_j$, which is only possible if the corresponding cluster connects two chains.

The two smallest clusters that connect two chains are shown in Fig. F.1 (a). They both have two links and their smallest contribution is of 4th order. Inspecting Eq. (F.9), we see that any term that is proportional to η_c also has a factor σ_i^β . Hence, the two clusters marked in red and blue will cancel out exactly. This will be true for any cluster that has a ‘mirrored’ cluster that connects to a site on the same chain but on a different sublattice. Hence, only clusters of nontrivial topology, i.e. those that have noncontractible loops, will contribute to the effective chain couplings. The smallest such cluster on the pyrochlore lattice is a hexagon and there are two inequivalent ways a hexagon can connect two chains, shown in Fig. 9.5 (a). Note that for the same reasons, there will be no nonvanishing contribution from perturbation theory, that is proportional to only a single chain variable η_i . Any such term must be proportional to a respective σ_i^β and has a mirrored term that gives the same contribution but with the opposite sign.

Evaluating the energy contributions of the two types of hexagons, we have to consider all possible distributions of $V_{ij}^{(1)}$ and $V_{ij}^{(2)}$. However, not all of those contribute to the chain coupling. Consider for example the contribution from the cluster depicted in Fig. F.1 (b), which is [using Eq. (F.9)]

$$\begin{aligned} \delta E^{(6)} \propto & (\cos \phi J_x + \eta_1 J_y) (\cos \phi J_x - \eta_1 J_y) \\ & (\cos \phi J_x + \eta_2 J_y) (\cos \phi J_x - \eta_2 J_y) (J_x - J_y)^2 \end{aligned} \quad (\text{F.12})$$

and thus the dependence on η_i drops out completely.

Keeping only those configurations which do depend on the chain configuration $\{\eta_i\}$, there are four inequivalent possibilities, shown in Fig. F.1 (c). Each of them have different possible sequences of intermediate states $\{\psi_i\}$, which must be enumerated to evaluate the denominator of Eq. (F.8). We show one possible such sequence, for the first configuration in Fig. F.1 (c), in Fig. F.1 (d), where the excited states are shown as dots on the sites.

Finally, the effective coupling J_{ch} is given as the sum of the four inequivalent contributions depicted in Fig. F.1 (c)

$$\begin{aligned} J_{\text{ch}} = & \frac{4(2S)^6}{4^6} \cos(\phi)^2 J_x^2 J_y^2 ((J_x - J_y)^2 (\Gamma_1 - \Gamma_2) \\ & + (J_x + J_y)^2 (\Gamma_3 - \Gamma_4)), \end{aligned} \quad (\text{F.13})$$

where we have defined factors Γ_i , stemming from the denominator of Eq. (F.8)

$$\Gamma_{\text{config}} := - \sum_{\{\psi_i\}} \frac{1}{(E_0 - E_{\psi_1}) \dots (E_0 - E_{\psi_{p-1}})} \quad (\text{F.14})$$

Note that since all factors in the denominator are negative and there are an odd number of factors for each closed loop on the pyrochlore lattice, $\Gamma_{\text{config}} > 0$.

Evaluating the Γ factors, one finds that $\Gamma_3 = \Gamma_4$ and hence

$$J_{\text{ch}}^z = - S J_z \cos(\phi)^2 \frac{J_x^2 J_y^2}{J_z^4} \frac{(J_x - J_y)^2}{J_z^2} \frac{1}{16} \Gamma_z(h), \quad (\text{F.15a})$$

$$\begin{aligned}
 & \text{Configuration 1: } = \text{const} + (2S)^6 \frac{4}{4^6} \cos^2 \phi^2 J_x^2 J_y^2 (-\cos \phi^2 J_x + J_y - \sin \phi^2 J_z)^2 \Gamma'_1 \eta_1 \eta_3 \\
 & \text{Configuration 2: } = \text{const} - (2S)^6 \frac{4}{4^6} \cos^2 \phi^2 J_x^2 J_y^2 (-\cos \phi^2 J_x + J_y - \sin \phi^2 J_z)^2 \Gamma'_2 \eta_1 \eta_3 \\
 & \text{Configuration 3: } = \text{const} + (2S)^6 \frac{4}{4^6} \cos^2 \phi^2 J_x^2 J_y^2 (-\cos \phi^2 J_x - J_y - \sin \phi^2 J_z)^2 \Gamma'_3 \eta_1 \eta_3 \\
 & \text{Configuration 4: } = \text{const} - (2S)^6 \frac{4}{4^6} \cos^2 \phi^2 J_x^2 J_y^2 (-\cos \phi^2 J_x - J_y - \sin \phi^2 J_z)^2 \Gamma'_4 \eta_1 \eta_3
 \end{aligned}$$

Figure F.2.: The four inequivalent configurations of perturbations on a hexagon contributing to J'_{ch} . We label the sites with reference to the sublattice \mathcal{L}_ν they belong to, consistent with the hexagon indicated in pink in Fig. 9.5 (a)

where $\Gamma_z(h) := S^5 J_z^5 (\Gamma_2 - \Gamma_1)$ is defined as a dimensionless factor controlling the asymptotic field dependence of the couplings and we have introduced a superscript ‘z’ to denote that the above result is true in the Chain_z phase.

Computing J'_{ch} follows exactly the same procedure. Since the sites on the hexagon, that is the leading-order contribution to J'_{ch} are on different sublattices, also configurations of perturbation on the hexagon that contribute will be different. The four inequivalent configurations that contribute to J'^z_{ch} , together with their contributions are shown in Fig. F.2. Summing those up yields

$$J'^z_{\text{ch}} = S J_z \cos^4 \phi^2 \frac{J_x^2 J_y^3}{J_z^5} (J_x/J_z + \tan \phi^2) \frac{2}{16} \Gamma'_z(h). \quad (\text{F.15b})$$

This completes the derivation of the results stated in the main text. For completeness, we give the Γ -factors explicitly:

$$\frac{\Gamma_z(h)}{(S J_z)^5} = \frac{2}{B_z^\beta (B_z^\alpha + B_z^\beta)^4} + \frac{2}{(B_z^\beta)^2 (B_z^\alpha + B_z^\beta)^3} + \frac{1}{(B_z^\beta)^3 (B_z^\alpha + B_z^\beta)^2}, \quad (\text{F.16a})$$

$$\frac{\Gamma'_z(h)}{(S J_z)^5} = \frac{4}{B_z^\beta (B_z^\alpha + B_z^\beta)^4}, \quad (\text{F.16b})$$

where B_z^α and B_z^β are the local fields on the α and β chains respectively, as defined in Eq. (F.11). Note that since the local fields B_z^α and B_z^β are always positive, so are Γ_z and Γ'_z

In the Chain_X phase, the ground state as defined in Eq. (F.5) will be slightly different

$$\hat{\boldsymbol{w}}_i = \begin{cases} (\sigma_i^\alpha \sin \phi, 0, \sigma_i^\alpha \cos \phi) & i \in \alpha \text{ chain} \\ \left(\sigma_i^\beta \eta_c, 0, 0 \right) & i \in \beta \text{ chain} \end{cases}, \quad (\text{F.17a})$$

and

$$\boldsymbol{c}_i = \frac{1}{2} \begin{cases} (\sigma_i^\alpha \cos \phi, i, -\sigma_i^\alpha \sin \phi) & i \in \alpha \text{ chain} \\ \left(0, i\sigma_i^\beta \eta_c, -1 \right) & i \in \beta \text{ chain} \end{cases}. \quad (\text{F.17b})$$

Note the sign change in the definition of \boldsymbol{c}_i on β chains that is necessary so the basis stays right-handed. Since \boldsymbol{J} is diagonal, the above change can be accounted for by simply changing $J_x \leftrightarrow J_z$ and $\cos \phi \leftrightarrow \sin \phi$ in Eq. (F.15) and hence

$$J_{\text{ch}}^x = -S J_x \sin(\phi)^2 \frac{J_z^2 J_y^2 (J_z - J_y)^2}{J_x^4 J_x^2} \frac{1}{16} \Gamma_x(h), \quad (\text{F.18a})$$

$$J_{\text{ch}}^{x'} = S J_x \sin(\phi)^4 \frac{J_x^2 J_y^3}{J_x^5} (J_z/J_x + \cotan \phi^2) \frac{2}{16} \Gamma'_x(h). \quad (\text{F.18b})$$

where $\Gamma_x^{(\prime)}$ is obtained from $\Gamma_z^{(\prime)}$ [Eq. (F.16)] by substituting B_z^α and B_z^β by B_x^α and B_x^β respectively.

F.3. Chain_Y Phase

We now turn to calculate the effective couplings in the Chain_Y phase, which is done mostly by the same procedure as used for the Chain_{X/Z} phases. As before, we start by defining a basis where the $\hat{\boldsymbol{w}}_i$ are given by the classical ground state configuration

$$\hat{\boldsymbol{w}}_i = \begin{cases} (\sigma_i^\alpha \sin \phi, 0, \sigma_i^\alpha \cos \phi) & i \in \alpha \text{ chain} \\ \left(0, \sigma_i^\beta \eta_c, 0 \right) & i \in \beta \text{ chain} \end{cases}, \quad (\text{F.19a})$$

and the \boldsymbol{c}_i complete the basis

$$\boldsymbol{c}_i = \frac{1}{2} \begin{cases} (\sigma_i^\alpha \cos \phi, i, -\sigma_i^\alpha \sin \phi) & i \in \alpha \text{ chain} \\ \left(i, 0, \sigma_i^\beta \eta_c \right) & i \in \beta \text{ chain} \end{cases}. \quad (\text{F.19b})$$

Substituting the above into Eq. (F.4) yields if i, j are both on a β chain

$$V_{ij}^{(1)} = -\frac{1}{4} (J_x + J_z) S_i^+ S_j^+ + h.c., \quad (\text{F.20a})$$

$$V_{ij}^{(2)} = \frac{1}{4} (J_x - J_z) S_i^+ S_j^- + h.c., \quad (\text{F.20b})$$

in contrast, if i, j are both on an α -chain

$$V_{ij}^{(1)} = -\frac{1}{4} (\cos \phi^2 J_x + J_y + \sin \phi^2 J_z) S_i^+ S_j^+ + h.c., \quad (\text{F.20c})$$

$$V_{ij}^{(2)} = -\frac{1}{4} (\cos \phi^2 J_x - J_y + \sin \phi^2 J_z) S_i^+ S_j^- + h.c., \quad (\text{F.20d})$$

and finally if i is on a β -chain and j is on an α -chain

$$V_{ij}^{(1)} = -\frac{1}{4}\sigma_j^\alpha \left(i \cos \phi J_x + \sigma_i^\beta \eta_c \sin \phi J_z \right) S_i^+ S_j^+ + h.c., \quad (\text{F.20e})$$

$$V_{ij}^{(2)} = -\frac{1}{4}\sigma_j^\alpha \left(i \cos \phi J_x + \sigma_i^\beta \eta_c \sin \phi J_z \right) S_i^+ S_j^- + h.c.. \quad (\text{F.20f})$$

Note that in the last case, the matrix elements are now complex.

F.3.1. Effective coupling J_{ch}^y

We now turn to identify the degeneracy-breaking perturbations. In the case of the $\text{Chain}_{Z/X}$ phases those are given by linked clusters of nontrivial topology. Note that the argument for this relied solely on the fact that any Ising chain variable η_c comes also with a factor σ_i^β , which is clearly still the case. Hence we can again restrict the discussion to the two kinds of hexagons as long as we are interested only in leading-order corrections.

As before, not all distributions of the perturbations $V_{ij}^{(1)}$ and $V_{ij}^{(2)}$ yield an energy correction that is proportional to the two Ising variables η_1, η_2 on the hexagon. For that, first consider the case of J_{ch}^y , which is generated by the hexagon marked in green in Fig. 9.5 (a). In this case, proportionality to a chain variable η_1 comes from a product of two perturbations connecting that chain to a polarized site, where the factor σ_i^β is different for both. Inspecting Eq. (F.20), we see that the two perturbations have to act differently on the two β chain sites in order for the resulting energy correction to be proportional to η_1 . Note that this rule already excludes the possibility of having $V_{ij}^{(1)}$ act on the β chain bonds [see Fig. F.3 (a)] and hence $J_{\text{ch}} \propto (J_x - J_z)^2$ in accordance with the LSWT result (see Fig. 9.7 in the main text). Fixing the perturbations acting on β chains to spin-flip hopping, there are four diagrams contributing to J_{ch}^y shown in Fig. F.3 (b). Summing up their contributions yields

$$J_{\text{ch}}^y = -\frac{1}{64} S J_y \sin(2\phi)^2 \frac{J_x^2 J_z^2 (J_x - J_z)^2}{J_y^6} \Gamma_y(h), \quad (\text{F.21})$$

$$\Gamma_y = S^5 J_y^5 (2\Gamma_2^Y + 2\Gamma_3^Y - 2\Gamma_1^Y - \Gamma_4^Y), \quad (\text{F.22})$$

where the Γ_i^Y are computed as before by enumerating all possible intermediate states and Γ_y is given by

$$\begin{aligned} \frac{\Gamma_y}{(S J_y)^5} &= \frac{2}{(B_y^\beta)^3 (B_y^\alpha + B_y^\beta)^2} - \frac{2}{(B_y^\alpha + 2B_y^\beta)^3 (B_y^\alpha + B_y^\beta)^2} \\ &+ \frac{4}{(B_y^\beta)^2 (B_y^\alpha + B_y^\beta)^3} - \frac{4}{(B_y^\alpha + 2B_y^\beta)^2 (B_y^\alpha + B_y^\beta)^3} \\ &+ \frac{2}{B_y^\beta (B_y^\alpha + B_y^\beta)^4} - \frac{2}{(B_y^\alpha + 2B_y^\beta)(B_y^\alpha + B_y^\beta)^4}, \end{aligned} \quad (\text{F.23})$$

where B_y^α and B_y^β are the local fields on the α and β chains respectively in the Chain_Y phase, obtained from Eq. (F.11) by replacing J_z by J_y . Note that since the local fields are always positive, it is clear that $\Gamma_y > 0$ and hence the effective chain coupling $J_{\text{ch}}^y < 0$ is ferromagnetic for all fields and exchange parameters.

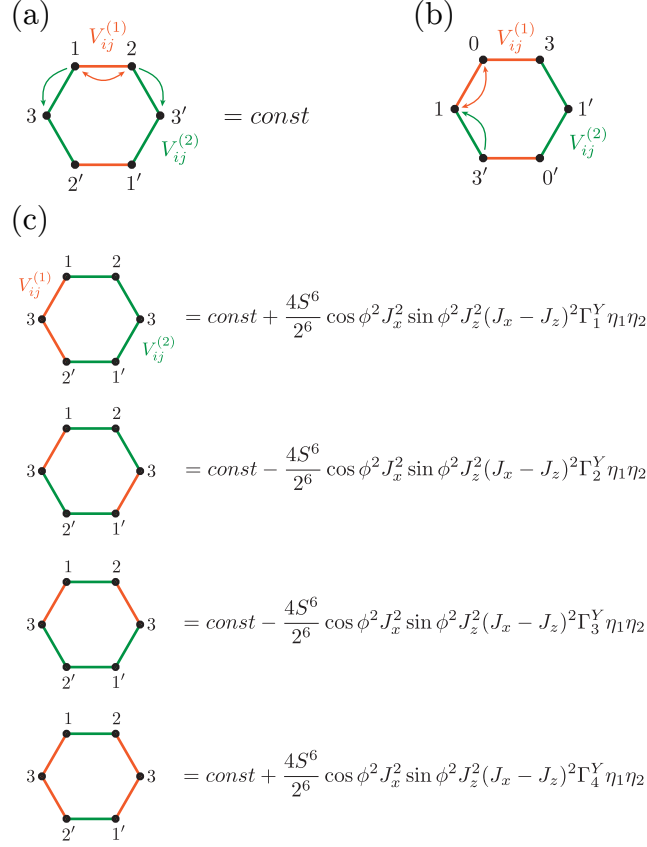


Figure F.3.: Diagrams arising in RSPT in the Chain_Y phase. (a) Example of a configuration that does not contribute, with one possible direction of spin flip hopping indicated. Whenever there is a double spin flip acting on the β chain (site 1, 2), the two perturbations on edge (1, 2) and (2, 3') will have to act in the same way (i.e. both with S^+ or both with S^-) on site 1 and 2 respectively, leading to a constant contribution. (b) The four perturbations contributing to J_{ch}^y . Note that the product of matrix elements does not depend on the order of application of the perturbations, so that the Γ factors can be computed as before. (c) A distribution of perturbations on the hexagon whose product of matrix elements yields a contribution to $J_{\text{ch}}^{y'}$. However, the diagram is not valid since there is no possible order in which, starting from the ground state, it yields the ground state also as a final state. For spin-1/2, there are actually no allowed diagrams using only $V_{ij}^{(1)}$, $V_{ij}^{(2)}$ and contributing to J'_{ch} , even at higher orders in RSPT.

F.3.2. Effective coupling $J_{\text{ch}}^{y'}$

We now turn to the second effective exchange coupling, $J_{\text{ch}}^{y'}$, which is generated by the hexagon marked in pink in Fig. 9.5 (a). On the level of RSPT diagrams, this amounts to a different distribution of polarized and unpolarized sites on the hexagon as indicated in Fig. F.3 (c). Turning to identify the degeneracy breaking perturbations, we straightforwardly arrive at a similar rule as in the case of J_{ch}^y . However, since the two perturbations connecting the site on the β chain now connect to the same site, and because σ_j^α appears just as an overall sign in the matrix elements [Eq. (F.20)], the two perturbations now have to act in the same way (i.e. either both with S^+ or both with S^-) on the site on the β chains as indicated in Fig. F.3 (c). However, this is not possible at sixth order since there can be no double occupation of excitations on any site. For the same reason, for $S = 1/2$ any diagram using only $V_{ij}^{(1)}$, $V_{ij}^{(2)}$ cannot yield a finite contribution to $J_{\text{ch}}^{y'}$, even at higher orders. Since in the Chain_Y phase $V_{ij}^{(3)} = 0$ if i on a β chain and j on an α chain (or the other way around), a finite contribution to $J_{\text{ch}}^{y'}$ can only arise from a perturbation including $V_{ij}^{(4)}$. Such a term occurs at seventh order in RSPT and competes with contributions from higher orders to J_{ch} , even when the sixth order contribution to J_{ch} vanishes.

Since in dipolar-octupolar pyrochlores the pseudospin always forms a doublet, we conclude that $J_{\text{ch}}^{y'}$ is strongly suppressed for these compounds. Taken together with the fact that $J_{\text{ch}}^y < 0$, we conclude that quantum fluctuations should mostly or even always select a ferromagnetic chain configuration in the Chain_Y phase. LSWT (see Fig. 9.7 of the main text) reproduces the overall dominance of ferromagnetic ordering at high fields, but predicts a zigzag ground state for some exchange parameters at low fields. We ascribe this to the fact that LSWT does not capture the hardcore nature of magnons at $S = 1/2$ since $V_{ij}^{(4)}$ corresponds to cubic magnon terms in spin wave theory.

Notwithstanding the ambiguity of the two methods at low fields, at higher fields both LSWT and RSPT calculations agree that ferromagnetic chain configurations will dominate the Chain_Y phase [see Fig. 9.7 (b) and (d)].

G. Spin-Wave Theory of Chain Phases in Dipolar-Octupolar Pyrochlores

In this appendix, we provide details about the linear spin wave theory calculations presented in the main text. We start from a system consisting of L unit cells U_l ($l = 0 \dots L - 1$), repeated along one dimension. Each unit cell has M sites at positions $l\mathbf{a} + \mathbf{r}_i$ ($i = 0 \dots M - 1$). The general nearest-neighbor exchange Hamiltonian of such a system reads

$$\begin{aligned} \mathcal{H} = & \frac{1}{2} \sum_l \sum_{i,j \in U_l} \mathbf{S}_i \mathcal{J}_{ij} \mathbf{S}_j + \sum_l \sum_{i \in U_l} \sum_{j \in U_{l+1}} \mathbf{S}_i \mathcal{K}_{ij} \mathbf{S}_j \\ & - \sum_l \sum_{i \in U_l} \mathbf{h}_i \cdot \mathbf{S}_i, \end{aligned} \quad (\text{G.1})$$

where \mathcal{J}_{ij} and \mathcal{K}_{ij} parametrize the exchange within and between unit cells respectively.

To derive the linear spin wave Hamiltonian, we first write the spins in a local basis $\mathbf{u}_i, \mathbf{v}_i, \mathbf{w}_i$, chosen such that in the classical ground state all spins align with \mathbf{w}_i

$$\mathbf{S}_i = S_i^u \mathbf{u}_i + S_i^v \mathbf{v}_i + S_i^w \mathbf{w}_i, \quad (\text{G.2})$$

and define raising and lowering operators via the (linear) Holstein-Primakoff transformation

$$S_i^u \approx \sqrt{S/2} (a_i + a_i^\dagger), \quad (\text{G.3a})$$

$$S_i^v \approx \sqrt{S/2} (a_i - a_i^\dagger), \quad (\text{G.3b})$$

$$S_i^w = S - a_i^\dagger a_i. \quad (\text{G.3c})$$

Then, we perform a one-dimensional Fourier transform on sublattice i

$$a_i^\dagger(\mathbf{q}) = \frac{1}{\sqrt{L}} \sum_l \exp(-i\mathbf{q} \cdot (l\mathbf{a} + \mathbf{r}_i)), \quad (\text{G.4})$$

where the momentum takes values $\mathbf{q} = 2\pi\mathbf{a}k/|\mathbf{a}|^2L$ ($k = 0 \dots L - 1$). Using the above, we arrive at

$$\mathcal{H} = E_0 + \mathcal{H}_1 + \mathcal{H}_2, \quad (\text{G.5})$$

where E_0 is the classical ground state energy, \mathcal{H}_1 contains terms linear in the operators a_i, a_i^\dagger and vanishes if $\mathbf{S}_i^{(0)} = S\hat{\mathbf{w}}_i$ is a classical ground state. Finally, \mathcal{H}_2 is

given by

$$\mathcal{H}_2 = \frac{1}{2} \sum_{\mathbf{q}} \mathbf{A}(\mathbf{q})^\dagger \mathcal{X}(\mathbf{q}) \mathbf{A}(\mathbf{q}), \quad (\text{G.6a})$$

$$\mathbf{A}^\dagger(\mathbf{q}) = (a_0^\dagger(\mathbf{q}), \dots, a_{M-1}^\dagger(\mathbf{q}), a_0(-\mathbf{q}), \dots), \quad (\text{G.6b})$$

$$\mathcal{X}(\mathbf{q}) = S \begin{pmatrix} X^{(11)}(\mathbf{q}) & X^{(12)}(\mathbf{q}) \\ X^{(21)}(\mathbf{q}) & X^{(22)}(\mathbf{q}) \end{pmatrix}, \quad (\text{G.6c})$$

$$\begin{aligned} X_{ij}^{(11)}(\mathbf{q}) = e^{-i\mathbf{q}\cdot\mathbf{r}_{ij}} & \left(\mathbf{c}_i \mathcal{J}_{ij} \mathbf{c}_j^* + e^{-i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_i \mathcal{K}_{ij} \mathbf{c}_j^* + e^{i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_j^* \mathcal{K}_{ji} \mathbf{c}_i \right. \\ & \left. - \delta_{ij} \left(\sum_{l=0\dots M-1} (\mathbf{w}_i \mathcal{J}_{il} \mathbf{w}_l + \mathbf{w}_i \mathcal{K}_{il} \mathbf{w}_l + \mathbf{w}_l \mathcal{K}_{li} \mathbf{w}_i) - \frac{1}{S} \mathbf{h}_i \cdot \mathbf{w}_i \right) \right), \end{aligned} \quad (\text{G.6d})$$

$$X_{ij}^{(12)}(\mathbf{q}) = e^{-i\mathbf{q}\cdot\mathbf{r}_{ij}} (\mathbf{c}_i \mathcal{J}_{ij} \mathbf{c}_j + e^{-i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_i \mathcal{K}_{ij} \mathbf{c}_j + e^{i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_j \mathcal{K}_{ji} \mathbf{c}_i), \quad (\text{G.6e})$$

$$X_{ij}^{(21)}(\mathbf{q}) = e^{-i\mathbf{q}\cdot\mathbf{r}_{ij}} (\mathbf{c}_i^* \mathcal{J}_{ij} \mathbf{c}_j^* + e^{-i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_i^* \mathcal{K}_{ij} \mathbf{c}_j^* + e^{i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_j^* \mathcal{K}_{ji} \mathbf{c}_i^*), \quad (\text{G.6f})$$

$$\begin{aligned} X_{ij}^{(22)}(\mathbf{q}) = e^{-i\mathbf{q}\cdot\mathbf{r}_{ij}} & \left(\mathbf{c}_i^* \mathcal{J}_{ij} \mathbf{c}_j + e^{-i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_i^* \mathcal{K}_{ij} \mathbf{c}_j + e^{i\mathbf{q}\cdot\mathbf{a}} \mathbf{c}_j \mathcal{K}_{ji} \mathbf{c}_i^* \right. \\ & \left. - \delta_{ij} \left(\sum_{l=0\dots M-1} (\mathbf{w}_i \mathcal{J}_{il} \mathbf{w}_l + \mathbf{w}_i \mathcal{K}_{il} \mathbf{w}_l + \mathbf{w}_l \mathcal{K}_{li} \mathbf{w}_i) - \frac{1}{S} \mathbf{h}_i \cdot \mathbf{w}_i \right) \right), \end{aligned} \quad (\text{G.6g})$$

where the $X^{(nm)}(\mathbf{q})$ are $M \times M$ matrices and $\mathbf{c}_i = \frac{1}{2}(\hat{\mathbf{u}}_i + i\hat{\mathbf{v}}_i)$. The sum goes over $\mathbf{q} = 2\pi\mathbf{a}k/|\mathbf{a}|^2L$ ($k = 0 \dots L-1$) and \mathbf{a} is the translation vector between neighboring unit cells and $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ is the difference vector between two sites within the unit same cell.

Eq. (G.6) is quadratic and thus can be diagonalized using a Bogoliubov transformation. Explicitly, the magnon dispersion is obtained by diagonalizing $\sigma\mathcal{X}(\mathbf{q})$ where

$$\sigma = \begin{pmatrix} I_{M \times M} & 0 \\ 0 & -I_{M \times M} \end{pmatrix}. \quad (\text{G.7})$$

The matrix $\sigma\mathcal{X}(\mathbf{q})$ has eigenvalues $\omega_\nu(\mathbf{q})$, $-\omega_\nu^*(\mathbf{q})$, $\nu = 0 \dots M-1$.

The magnon spectrum $\omega_\nu(\mathbf{q})$ depends on the classical ground state $\{\mathbf{S}_i^{(0)}\}$ used in the expansion above. The ground state selection by quantum fluctuations can then be quantified using the zero point energy

$$\mathcal{E}_0(\{\mathbf{S}_i^{(0)}\}) = \lim_{L \rightarrow \infty} \frac{1}{L} \langle 0 | H_2 | 0 \rangle \quad (\text{G.8})$$

$$= \frac{|\mathbf{a}|}{4\pi} \sum_{\nu} \int_0^{2\pi/|\mathbf{a}|} dk \omega_\nu(k\mathbf{a}/|\mathbf{a}|). \quad (\text{G.9})$$

G.1. Details of the Numerical Calculations

Here, we provide some details of the computation of the numerical phase diagram in Fig. 9.7. For a discussion of the different phases, we refer to the main text [Sec. 9.2.2].

As mentioned already in the main text, the LSWT theory calculations are performed on a semi-infinite “tube” of the pyrochlore lattice, that is an effective $W \times W$ triangular lattice of β chains. The infinite direction is thus $[1\bar{1}0]$ [$\mathbf{a} = (2, -2, 0)$] and the cluster has $M = 4W^2$ sites per unit cell. The zero point energy is then given by

$$\mathcal{E}_0(\{\mathcal{S}_i^{(0)}\}) = \frac{1}{\pi} \sum_{\nu} \int_0^{\pi/2} dk \omega_{\nu}((k, -k, 0)), \quad (\text{G.10})$$

where we used that the dispersion is symmetric around $k = 0$.

For a fixed set of exchange couplings \mathbf{J} and field \mathbf{h} , we diagonalize the resulting $2M \times 2M$ matrix numerically for 50 values of \mathbf{q} and compute the integral of the dispersion numerically using the standard Simpson’s quadrature rule. Since relative differences between different zero point energies are very small, we verified that using 50 points, the integral is converged up to double precision. For $W = 8$, computing the zero point energy of a single classical ground state for fixed parameters takes about 10 seconds.

The number of distinct configurations of the Ising variables $\{\eta_i\}$ grows exponentially with the system size W . The last system size for which it is feasible to enumerate all possible configurations is $W = 4$, for which there are $2^{W^2} = 65536$. This can be reduced to 674 inequivalent chain configurations which are not related by symmetries. Computing the zero point energies [Eq. (G.10)] for all these states in linear spin wave theory yields the phase boundary between the ferromagnetic and the zigzag state, as well as the gap Δ between the two. However, at this system size, LSWT is not able to correctly capture the splitting between zigzag states. This is because the unit cell of the two chosen ground states, the kinked and the striped states is of size $W = 4$. If we would simulate the effective triangular lattice Ising model directly, a system of the size of that unit cell would be sufficient to capture the selection process.

However, since we actually consider the whole pyrochlore lattice and wish to capture the fluctuations which generate the interactions between the chains, we need a cluster which is at least twice as large to suppress the influence of virtual processes which cross the entire cluster via the periodic boundary conditions.

In order to obtain the fully resolved phase diagram shown in Fig. 9.7, we use a system with size $W = 8$. For each set of parameters, we compute the zero point energy of all possible ground states of the effective model, that are the ferromagnet, the staggered state, and all inequivalent zigzag states. The latter is possible because the total number of zigzag states grows exponentially only in the linear system size, hence there are $2^W = 256$ zigzag states in total. The number can be further reduced to 12 zigzag states not related by symmetry. To summarize, there are 14 possible inequivalent ground states of the effective model for $W = 8$.

To benchmark the validity of the effective triangular lattice Ising model we calculate the zero-point energy of 1000 random states. We show the results of this benchmarking in Fig. G.1. In (a) we compare the zero point energies obtained by LSWT to those computed from an effective model with effective couplings fitted using three states [Eq. (9.20)]. In (b), we compare the effective chain couplings obtained from 6th order RSPT ($J_{\text{ch}}^{(l)(6)}$) to those computed from LSWT ($J_{\text{ch}}^{(l)\text{SW}}$).

Finally, to assess finite size effects we also simulate a system with $W = 16$ and verify that the numerical values of Δ and B are independent of the size.

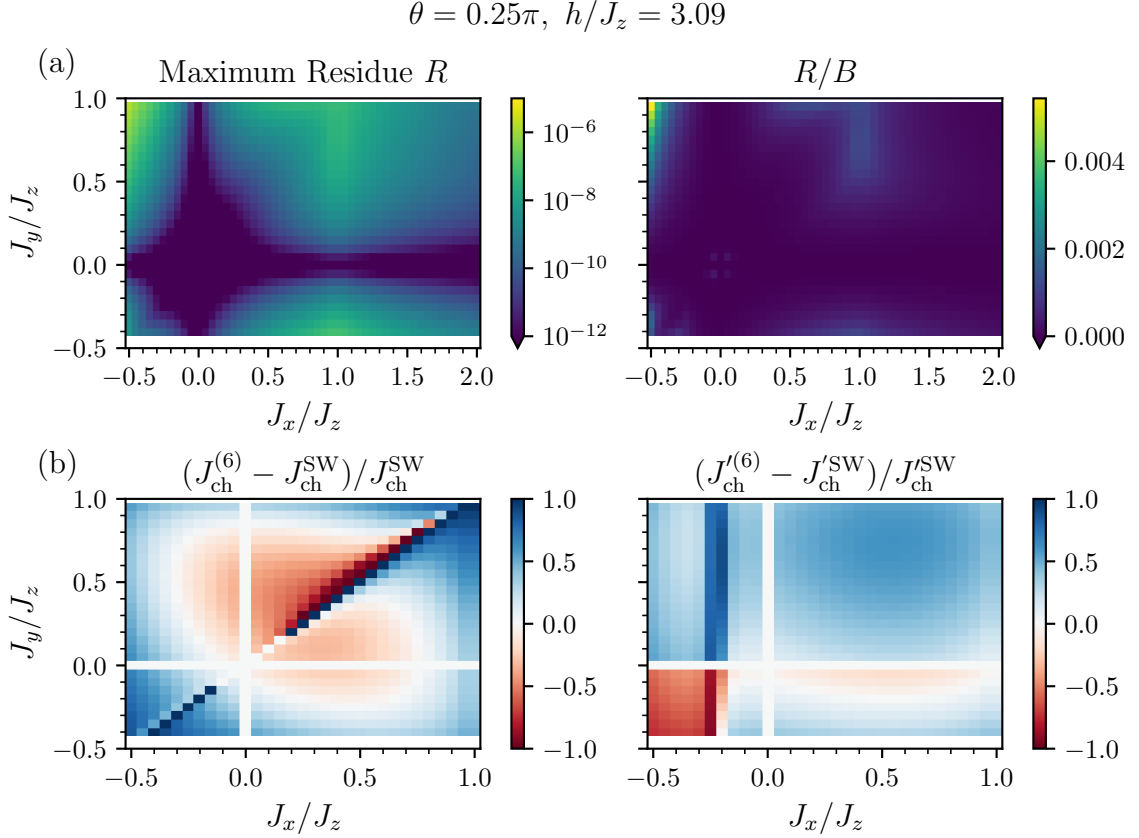


Figure G.1.: Benchmarking the effective model using linear spin wave theory (LSWT) using the 14 ground states of the effective model and 1000 random chain configurations. (a) Maximum residue R between a zero point energy (per site) predicted by the effective model and computed by LSWT. We give its absolute value as well as divided by the total bandwidth of zero point energies B , which is a scale given by the effective nearest-neighbor chain couplings. (b) Comparison between the nearest-neighbor effective chain couplings as calculated from 6th real-space perturbation theory ($J_{\text{ch}}^{(6)}$) and as obtained from LSWT using Eq. (9.20) ($J_{\text{ch}}^{\text{SW}}$).

G.2. Survival of Full Classical Degeneracy when $J_x = 0$ or $J_y = 0$ in the Dipolar Chain Phases

In this appendix, we prove the survival of the full classical degeneracy, within LSWT, in the singular cases of $J_x = 0$ or $J_y = 0$. To this end, we relate the flipping of single β chain in the system to a unitary transformation on the matrix $\sigma\mathcal{X}(\mathbf{q})$. This already implies that the spectrum and hence the zero point energy of the system does not depend on the chain configuration $\{\eta_i\}$ and hence the full classical degeneracy is preserved even in the presence of quantum zero-point fluctuations.

In the dipolar chain phase i.e. in either Chain_X or Chain_Z , all classical ground states are periodic in the $[1\bar{1}0]$ direction with a unit cell of four sites per β chain in the system. Two sites of those are on the β chain and the other two are on α chains. We now consider the case of a general (nonzero) matrix element of $\mathcal{X}(\mathbf{q})$ as given in Eq. (G.6). Flipping a single chain variable $\eta_c \rightarrow -\eta_c$ amounts to a change in a maximum of 18 matrix elements in each of the four submatrices $X^{(mn)}$ of $\mathcal{X}(\mathbf{q})$. These matrix elements can take two forms.

In the following, we assume for simplicity to be in the Chain_Z phase, but the same argument applies in the Chain_X phase. First, if i, j are on the same β chain, the matrix elements are given by

$$X_{ij}^{(11)}(\mathbf{q}) = X_{ij}^{(22)}(\mathbf{q}) = \cos(\pi k/L) (J_x - J_y), \quad (\text{G.11a})$$

$$X_{ij}^{(12)}(\mathbf{q}) = X_{ij}^{(21)}(\mathbf{q}) = \cos(\pi k/L) (J_x + J_y), \quad (\text{G.11b})$$

where we used the definition of \mathbf{c}_i from Eq. (F.5). The above is independent of the Ising variable η_c of the chain and hence there are only 16 matrix elements left that might change as the chain is flipped. The other sixteen matrix elements are all of the form

$$X_{ij}^{(11)}(\mathbf{q}) = X_{ij}^{(22)}(\mathbf{q}) = \frac{1}{2} \exp(\pm i\mathbf{q} \cdot \mathbf{a}/4) \left(\sigma_j^\alpha J_x \cos \phi + \sigma_i^\beta \eta_c J_y \right), \quad (\text{G.12a})$$

$$X_{ij}^{(12)}(\mathbf{q}) = X_{ij}^{(21)}(\mathbf{q}) = -\frac{1}{2} \exp(\pm i\mathbf{q} \cdot \mathbf{a}/4) \left(\sigma_j^\alpha J_x \cos \phi - \sigma_i^\beta \eta_c J_y \right), \quad (\text{G.12b})$$

where i is a site on a β chain, j is a site on an α chain and $\sigma_i^\alpha, \sigma_j^\beta$ are defined in Eq. (F.6). The sign of the site-dependent phase $\pm i\mathbf{q} \cdot \mathbf{a}/4$ depends on the specific sites i and j , but this explicit dependence does not matter for the argument presented here. For $J_y = 0$, it is clear that the spin wave matrix $\mathcal{X}(\mathbf{q})$ does not depend on the Ising variable η_c and hence there is no ground state selection. For $J_x = 0$ the argument is less obvious but the same is still true. In that case, all matrix elements that change when $\eta_c \rightarrow -\eta_c$ are proportional to η_c . Thus, flipping the chain is equivalent to a simple unitary transformation on $\sigma\mathcal{X}(\mathbf{q})$ with the transformation matrix given by

$$U = \text{diag}(1, \dots, -1, -1, \dots, -1, -1, \dots, 1), \quad (\text{G.13})$$

where the -1 entries appear at the positions of the two sites on the flipped β chain.

G.3. Effect of Finite Temperature

In the following, we study the effect of finite temperature of ground state selection. To this end, we first consider classical spin wave theory, that is the low-temperature

expansion of the classical free energy around different ground states and show that the first order correction is independent of the ground state. Second, we consider the free energy of the magnons at finite temperature. We show that this leads to the same ground state selection for all temperatures. For temperatures below a crossover temperature T_{co} , set by the spin wave gap, the strength of ground state selection is close to the zero temperature value. For $T > T_{co}$ the difference in free energy between different chain configurations is heavily suppressed. For parameters of interest (close to those estimated for $\text{Nd}_2\text{Zr}_2\text{O}_7$) this crossover temperature is given by $T_{co} \approx 0.2J_z$ (see Fig. G.2).

G.3.1. Low temperature expansion

For the low temperature expansion, we choose as a starting point the same general one-dimensional Hamiltonian as for linear spin wave theory before. That is a system consisting of L unit cells U_l ($l = 0 \dots L - 1$), repeated along one dimension. Each unit cell has M sites at positions $l\mathbf{a} + \mathbf{r}_i$ ($i = 0 \dots M - 1$) and the Hamiltonian reads

$$\mathcal{H} = \frac{1}{2} \sum_l \sum_{i,j \in U_l} \mathbf{S}_i \mathcal{J}_{ij} \mathbf{S}_j + \sum_l \sum_{i \in U_l} \sum_{j \in U_{l+1}} \mathbf{S}_i \mathcal{K}_{ij} \mathbf{S}_j - \sum_l \sum_{i \in U_l} \mathbf{h}_i \cdot \mathbf{S}_i, \quad (\text{G.14})$$

We again proceed by writing the spin variables in a local basis $\mathbf{u}_i, \mathbf{v}_i, \mathbf{w}_i$ chosen such that in the (classical) ground state all spins align with \mathbf{w}_i . We can hence parametrize fluctuations around the ground state as

$$\begin{aligned} \mathbf{S}_i &= \sqrt{S} \delta u_i \mathbf{u}_i + \sqrt{S} \delta v_i \mathbf{v}_i + \sqrt{S^2 - S(\delta u_i)^2 - S(\delta v_i)^2} \mathbf{w}_i \\ &\approx \sqrt{S} \delta u_i \mathbf{u}_i + \sqrt{S} \delta v_i \mathbf{v}_i + \left(S - \frac{1}{2}(\delta u_i)^2 - \frac{1}{2}(\delta v_i)^2 \right) \mathbf{w}_i \end{aligned} \quad (\text{G.15})$$

substituting the above into Eq. (G.14) up to quadratic order in the fluctuations $\delta u_i, \delta v_i$ yields

$$\mathcal{H} = E_0 + \mathcal{H}_1 + \mathcal{H}_2 \quad (\text{G.16a})$$

where E_0 is the ground state energy, \mathcal{H}_1 is linear in the fluctuations $\delta u_i, \delta v_i$ and vanishes if $\mathbf{S}_i^{(0)} = S\mathbf{w}_i$ is a classical ground state, and \mathcal{H}_2 is given by

$$\mathcal{H}_2 = \frac{1}{2} \sum_{\mathbf{q}} \tilde{\mathbf{u}}(-\mathbf{q})^T \mathfrak{X}(\mathbf{q}) \tilde{\mathbf{u}}(\mathbf{q}), \quad (\text{G.16b})$$

$$\tilde{\mathbf{u}}(\mathbf{q})^T = (\delta u_1(q), \delta u_2(q), \dots, \delta u_M(q), \delta v_1(q), \dots). \quad (\text{G.16c})$$

Here, the matrix $\mathfrak{X}(\mathbf{q})$ is related to the spin wave Hamiltonian $\mathcal{X}(\mathbf{q})$ Eq. (G.6) derived in App. G by a unitary transformation

$$\mathfrak{X}(\mathbf{q}) = F^\dagger \mathcal{X}(\mathbf{q}) F \quad (\text{G.17a})$$

$$F = \frac{1}{\sqrt{2}} \begin{pmatrix} \mathbb{1} & i\mathbb{1} \\ \mathbb{1} & -i\mathbb{1} \end{pmatrix} \quad (\text{G.17b})$$

The classical partition function can be calculated explicitly

$$\begin{aligned} Z &= \text{tr} e^{-\beta \mathcal{H}} = (2\pi)^{-ML} \int \left(\prod_{\nu=1}^{2M} d\tilde{u}_\nu \right) e^{-\beta E_0} e^{-\beta \mathcal{H}_2} \\ &= e^{-\beta E_0} \prod_{\mathbf{q}} (\det \beta \mathfrak{X}(\mathbf{q}))^{\frac{1}{2}} \end{aligned} \quad (\text{G.18})$$




Figure G.2.: Order-by-disorder as a function of temperature for exchange parameters as estimated for $\text{Nd}_2\text{Zr}_2\text{O}_7$ (see Eq. (9.21)). The plot shows the gap Δ between the ferromagnetic chain configuration and the zigzag band. For low temperatures, it approaches the gap in zero point energies while for high temperatures, the gap vanishes.

and the free energy up to linear order in temperature by

$$\begin{aligned} F &= -T \log Z \\ &= E_0 + \frac{T}{2} \sum_{\mathbf{q}} \log \det \mathfrak{X}(\mathbf{q}) - TML \log T. \end{aligned} \quad (\text{G.19})$$

Now, note that flipping a chain in all cases amounts to flipping \mathbf{w}_i on two sites. To keep the basis right handed, one then also has to flip either \mathbf{u}_i or \mathbf{v}_i on those sites. In the expression for the partition function (G.18) however, assuming we flipped \mathbf{u}_i , one can revert this change by the transformation of variables $\delta u_i \rightarrow -\delta u_i$. Since the sign of an even number of variables is flipped, this leaves the phase space volume element $\left(\prod_{\nu=1}^{2M} d\tilde{u}_\nu\right)$ invariant. Hence, the partition function Z as well as the free energy F are invariant under flipping any single chain variable η_c up to second order in deviations from the ground state configuration. Generally the above argument implies that for a set of (locally) collinear ground states as considered here in all of the chain phases, there is no order-by-disorder when taking the purely classical limit.

G.3.2. Free energy of magnons at finite temperature

The second approach treating the finite temperature is considering the gas of non-interacting magnons. For that, we start from the spin wave Hamiltonian in Eq. (G.6), which after a Bogoliubov transformation takes the form

$$\mathcal{H}_2 = \sum_{\mathbf{q}, \nu} \omega_\nu(\mathbf{q}) \left(n_\nu(\mathbf{q}) + \frac{1}{2} \right) \quad (\text{G.20})$$

where $n_\nu(\mathbf{q}) = \eta_\nu^\dagger(\mathbf{q})\eta_\nu(\mathbf{q})$ is the magnon number operator and the magnon dispersion $\omega_\nu(\mathbf{q})$ is given by the eigenvalues of $\sigma\mathcal{X}(\mathbf{q})$. Note that these eigenvalue come in pairs $\pm\omega_\nu(\mathbf{q})$ (see App. G for details).

The free energy of this system at temperature T is then

$$F_2 = -T \log(\text{tr } e^{-\beta\mathcal{H}_2}) \quad (\text{G.21})$$

$$= \sum_{\mathbf{q},\nu} T \log(1 - e^{-\beta\omega_\nu(\mathbf{q})}) + \frac{1}{2}\omega_\nu(\mathbf{q}) \quad (\text{G.22})$$

where the second term in the sum yields the zero point energy and the first term is a finite-temperature correction. For zero temperature, the above is equal to the zero point energy while for large temperature it becomes independent of the classical ground state configuration $\{\mathbf{S}_j^{(0)}\}$. Particularly, for all temperatures the same ground state is favored by fluctuation of the free Bose gas but the gap in free energy vanishes as $T \rightarrow \infty$. As an example, we show the free energy gap Δ between the ferromagnetic chain configuration and the zigzag band as a function of temperature, for exchange parameters as estimated for $\text{Nd}_2\text{Zr}_2\text{O}_7$, in Fig. G.2.

The vanishing order-by-disorder in the high temperature limit can be understood by expanding the exponential inside the logarithm in Eq. (G.22)

$$\begin{aligned} F_2 &= \sum_{\mathbf{q},\nu} T \log\left(\beta\omega_\nu(\mathbf{q}) - \frac{1}{2}(\beta\omega_\nu(\mathbf{q}))^2 + \mathcal{O}((\beta\omega_\nu(\mathbf{q}))^3)\right) + \frac{1}{2}\omega_\nu(\mathbf{q}) \\ &= \sum_{\mathbf{q},\nu} T \log[\omega_\nu(\mathbf{q})] - T \log(T) + \mathcal{O}(\beta\omega_\nu(\mathbf{q})^2). \end{aligned} \quad (\text{G.23})$$

Note that the Eigenvalues of $\sigma\mathcal{X}(\mathbf{q})$ are exactly the pairs $\pm\omega_\nu(\mathbf{q})$ and the number of eigenvalues M is even such that $\det \sigma\mathcal{X}(\mathbf{q})$ is positive. Hence, we can write the leading order term of F_2 as

$$T \sum_{\mathbf{q},\nu} \log[\omega_\nu(\mathbf{q})] = \frac{T}{2} \sum_{\mathbf{q}} \log \det \sigma\mathcal{X}(\mathbf{q}). \quad (\text{G.24})$$

Now, using that $\det \sigma = 1$ and that the quantum spin wave matrix $\mathcal{X}(\mathbf{q})$ is related to its classical counterpart $\mathfrak{X}(\mathbf{q})$ (defined in Eq. (G.16)) by a unitary transformation [Eq. (G.17)], we arrive at

$$T \sum_{\mathbf{q},\nu} \log[\omega_\nu(\mathbf{q})] = \frac{T}{2} \sum_{\mathbf{q}} \log \det \mathfrak{X}(\mathbf{q}). \quad (\text{G.25})$$

At infinite temperature, we hence recover the purely classical result where, as discussed in the previous section of this appendix, there is no order by disorder.

We note that the discussion of finite-temperature OBD as presented here is not comprehensive as it does not incorporate the possible presence of domain walls at finite temperature. However, since in the presence of anisotropies domain wall excitations, just as spin wave excitations, are gapped, we expect our results to be robust in the low-temperature limit.

Bibliography

- [1] Heike Kamerlingh Onnes. “Further experiments with liquid helium. C. On the change of electric resistance of pure metals at very low temperatures etc. IV. The resistance of pure mercury at helium temperatures”. In: *Proceedings of the Section of Sciences* 13 (), pp. 1274–1276.
- [2] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. “Theory of Superconductivity”. In: *Phys. Rev.* 108 (5 Dec. 1957), pp. 1175–1204. DOI: 10.1103/PhysRev.108.1175.
- [3] K. v. Klitzing, G. Dorda, and M. Pepper. “New Method for High-Accuracy Determination of the Fine-Structure Constant Based on Quantized Hall Resistance”. In: *Phys. Rev. Lett.* 45 (6 Aug. 1980), pp. 494–497. DOI: 10.1103/PhysRevLett.45.494.
- [4] D. C. Tsui, H. L. Stormer, and A. C. Gossard. “Two-Dimensional Magnetotransport in the Extreme Quantum Limit”. In: *Phys. Rev. Lett.* 48 (22 May 1982), pp. 1559–1562. DOI: 10.1103/PhysRevLett.48.1559.
- [5] J. G. Bednorz and K. A. Müller. “Possible highT_c superconductivity in the Ba-La-Cu-O system”. In: *Zeitschrift für Physik B Condensed Matter* 64.2 (June 1986), pp. 189–193. ISSN: 1431-584X. DOI: 10.1007/BF01303701.
- [6] R. B. Laughlin. “Anomalous Quantum Hall Effect: An Incompressible Quantum Fluid with Fractionally Charged Excitations”. In: *Phys. Rev. Lett.* 50 (18 May 1983), pp. 1395–1398. DOI: 10.1103/PhysRevLett.50.1395.
- [7] Horst L. Stormer, Daniel C. Tsui, and Arthur C. Gossard. “The fractional quantum Hall effect”. In: *Rev. Mod. Phys.* 71 (2 Mar. 1999), S298–S305. DOI: 10.1103/RevModPhys.71.S298.
- [8] Walter. Kauzmann. “The Nature of the Glassy State and the Behavior of Liquids at Low Temperatures.” In: *Chemical Reviews* 43.2 (Oct. 1948), pp. 219–256. ISSN: 0009-2665. DOI: 10.1021/cr60135a002.
- [9] Frank H. Stillinger. “Supercooled liquids, glass transitions, and the Kauzmann paradox”. In: *The Journal of Chemical Physics* 88.12 (June 1988), pp. 7818–7825. ISSN: 0021-9606. DOI: 10.1063/1.454295.
- [10] Ludovic Berthier and Giulio Biroli. “Theoretical perspective on the glass transition and amorphous materials”. In: *Rev. Mod. Phys.* 83 (2 June 2011), pp. 587–645. DOI: 10.1103/RevModPhys.83.587.
- [11] P. W. Anderson. “More Is Different”. In: *Science* 177.4047 (1972), pp. 393–396. DOI: 10.1126/science.177.4047.393.
- [12] P.W. Anderson. “Resonating valence bonds: A new kind of insulator?” In: *Materials Research Bulletin* 8.2 (1973), pp. 153–160. ISSN: 0025-5408. DOI: [https://doi.org/10.1016/0025-5408\(73\)90167-0](https://doi.org/10.1016/0025-5408(73)90167-0).

- [13] F. D. M. Haldane. “Model for a Quantum Hall Effect without Landau Levels: Condensed-Matter Realization of the ”Parity Anomaly””. In: *Phys. Rev. Lett.* 61 (18 Oct. 1988), pp. 2015–2018. DOI: 10.1103/PhysRevLett.61.2015.
- [14] A.Yu. Kitaev. “Fault-tolerant quantum computation by anyons”. In: *Annals of Physics* 303.1 (2003), pp. 2–30. ISSN: 0003-4916. DOI: [https://doi.org/10.1016/S0003-4916\(02\)00018-0](https://doi.org/10.1016/S0003-4916(02)00018-0).
- [15] Claudio Chamon. “Quantum Glassiness in Strongly Correlated Clean Systems: An Example of Topological Overprotection”. In: *Phys. Rev. Lett.* 94 (4 Jan. 2005), p. 040402. DOI: 10.1103/PhysRevLett.94.040402.
- [16] C. L. Kane and E. J. Mele. “ Z_2 Topological Order and the Quantum Spin Hall Effect”. In: *Phys. Rev. Lett.* 95 (14 Sept. 2005), p. 146802. DOI: 10.1103/PhysRevLett.95.146802.
- [17] Alexei Kitaev. “Anyons in an exactly solved model and beyond”. In: *Annals of Physics* 321.1 (2006). January Special Issue, pp. 2–111. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2005.10.005>.
- [18] Vedika Khemani, Achilleas Lazarides, et al. “Phase Structure of Driven Quantum Systems”. In: *Phys. Rev. Lett.* 116 (25 June 2016), p. 250401. DOI: 10.1103/PhysRevLett.116.250401.
- [19] Daan Frenkel. “Order through entropy”. In: *Nature Materials* 14.1 (Jan. 2015), pp. 9–12. ISSN: 1476-4660. DOI: 10.1038/nmat4178.
- [20] Lars Onsager. “The Effects of Shape on the Interaction of colloidal particles”. In: *Annals of the New York Academy of Sciences* 51.4 (1949), pp. 627–659. DOI: <https://doi.org/10.1111/j.1749-6632.1949.tb27296.x>.
- [21] J Villain, R Bidaux, et al. “Order as an effect of disorder”. In: *Journal de Physique* 41.11 (1980), 1263–1272. ISSN: 0302-0738. DOI: 10.1051/jphys:0198000410110126300.
- [22] *Oxford Dictionary of Science*. 6th ed. Oxford University Press.
- [23] G. Toulouse. “Theory of the frustration effect in spin glasses: I”. In: *Communications on Physics* (2).
- [24] A P Ramirez. “Strongly Geometrically Frustrated Magnets”. In: *Annual Review of Materials Science* 24.1 (1994), pp. 453–480. DOI: 10.1146/annurev.ms.24.080194.002321.
- [25] John T. Chalker. “Spin liquids and frustrated magnetism”. In: *Topological Aspects of Condensed Matter Physics: Lecture Notes of the Les Houches Summer School: Volume 103, August 2014*. Oxford University Press, Jan. 2017. ISBN: 9780198785781. DOI: 10.1093/acprof:oso/9780198785781.003.0003.
- [26] R. Moessner and J. T. Chalker. “Properties of a Classical Spin Liquid: The Heisenberg Pyrochlore Antiferromagnet”. In: *Phys. Rev. Lett.* 80 (13 Mar. 1998), pp. 2929–2932. DOI: 10.1103/PhysRevLett.80.2929.
- [27] R. Moessner and J. T. Chalker. “Low-temperature properties of classical geometrically frustrated antiferromagnets”. In: *Phys. Rev. B* 58 (18 Nov. 1998), pp. 12049–12062. DOI: 10.1103/PhysRevB.58.12049.

- [28] J. T. Chalker, P. C. W. Holdsworth, and E. F. Shender. “Hidden order in a frustrated system: Properties of the Heisenberg Kagomé antiferromagnet”. In: *Phys. Rev. Lett.* 68 (6 Feb. 1992), pp. 855–858. DOI: 10.1103/PhysRevLett.68.855.
- [29] A. B. Harris, C. Kallin, and A. J. Berlinsky. “Possible Néel orderings of the Kagomé antiferromagnet”. In: *Phys. Rev. B* 45 (6 Feb. 1992), pp. 2899–2919. DOI: 10.1103/PhysRevB.45.2899.
- [30] I. Ritchey, P. Chandra, and P. Coleman. “Spin folding in the two-dimensional Heisenberg kagomé antiferromagnet”. In: *Phys. Rev. B* 47 (22 June 1993), pp. 15342–15345. DOI: 10.1103/PhysRevB.47.15342.
- [31] David A. Huse and Andrew D. Rutenberg. “Classical antiferromagnets on the Kagomé lattice”. In: *Phys. Rev. B* 45 (13 Apr. 1992), pp. 7536–7539. DOI: 10.1103/PhysRevB.45.7536.
- [32] M. E. Zhitomirsky. “Octupolar ordering of classical kagome antiferromagnets in two and three dimensions”. In: *Phys. Rev. B* 78 (9 Sept. 2008), p. 094423. DOI: 10.1103/PhysRevB.78.094423.
- [33] Gia-Wei Chern and R. Moessner. “Dipolar Order by Disorder in the Classical Heisenberg Antiferromagnet on the Kagome Lattice”. In: *Phys. Rev. Lett.* 110 (7 Feb. 2013), p. 077201. DOI: 10.1103/PhysRevLett.110.077201.
- [34] P. C. Guruciaga, M. Tarzia, et al. “Field-Tuned Order by Disorder in Frustrated Ising Magnets with Antiferromagnetic Interactions”. In: *Phys. Rev. Lett.* 117 (16 Oct. 2016), p. 167203. DOI: 10.1103/PhysRevLett.117.167203.
- [35] Patrik Fazekas. *Lecture Notes on Electron Correlation and Magnetism (Modern Condensed Matter Physics)*. English. Hardcover. World Scientific Publishing Company, Jan. 1999, p. 794. ISBN: 978-9810224745.
- [36] J. P. C. Ruff, J. P. Clancy, et al. “Spin Waves and Quantum Criticality in the Frustrated XY Pyrochlore Antiferromagnet $\text{Er}_2\text{Ti}_2\text{O}_7$ ”. In: *Phys. Rev. Lett.* 101 (14 Oct. 2008), p. 147205. DOI: 10.1103/PhysRevLett.101.147205.
- [37] M. E. Zhitomirsky, M. V. Gvozdikova, et al. “Quantum Order by Disorder and Accidental Soft Mode in $\text{Er}_2\text{Ti}_2\text{O}_7$ ”. In: *Phys. Rev. Lett.* 109 (7 Aug. 2012), p. 077204. DOI: 10.1103/PhysRevLett.109.077204.
- [38] Lucile Savary, Kate A. Ross, et al. “Order by Quantum Disorder in $\text{Er}_2\text{Ti}_2\text{O}_7$ ”. In: *Phys. Rev. Lett.* 109 (16 Oct. 2012), p. 167201. DOI: 10.1103/PhysRevLett.109.167201.
- [39] John Mauro. *Topological constraint theory of glass*. English (US). May 2011.
- [40] Mathieu Bauchy. “Topological Constraint Theory and Rigidity of Glasses”. In: *21st Century Nanoscience – A Handbook*. Ed. by K.D. Sattler. 1st ed. CRC Press, 2019. Chap. 13. DOI: 10.1201/9780367333003.
- [41] S F Edwards and P W Anderson. “Theory of spin glasses”. In: *Journal of Physics F: Metal Physics* 5.5 (May 1975), pp. 965–974. ISSN: 0305-4608. DOI: 10.1088/0305-4608/5/5/017.
- [42] K. Binder and A. P. Young. “Spin glasses: Experimental facts, theoretical concepts, and open questions”. In: *Rev. Mod. Phys.* 58 (4 Oct. 1986), pp. 801–976. DOI: 10.1103/RevModPhys.58.801.

- [43] J D Bryngelson and P G Wolynes. “Spin glasses and the statistical mechanics of protein folding.” In: *Proceedings of the National Academy of Sciences* 84.21 (1987), pp. 7524–7528. DOI: 10.1073/pnas.84.21.7524.
- [44] Marc Mézard and Riccardo Zecchina. “Random K -satisfiability problem: From an analytic solution to an efficient algorithm”. In: *Phys. Rev. E* 66 (5 Nov. 2002), p. 056126. DOI: 10.1103/PhysRevE.66.056126.
- [45] Pablo G. Debenedetti and Frank H. Stillinger. “Supercooled liquids and the glass transition”. In: *Nature* 410.6825 (Mar. 2001), pp. 259–267. ISSN: 1476-4687. DOI: 10.1038/35065704.
- [46] R. Brüning and K. Samwer. “Glass transition on long time scales”. In: *Phys. Rev. B* 46 (18 Nov. 1992), pp. 11318–11322. DOI: 10.1103/PhysRevB.46.11318.
- [47] M. D. Ediger, C. A. Angell, and Sidney R. Nagel. “Supercooled Liquids and Glasses”. In: *The Journal of Physical Chemistry* 100.31 (Jan. 1996), pp. 13200–13212. ISSN: 0022-3654. DOI: 10.1021/jp953538d.
- [48] H. Vogel. “Das Temperaturabhaengigkeitsgesetz der Viskositaet von Fluessigkeiten”. In: *Physikalische Zeitschrift* 22 (1921).
- [49] G. Tammann and W. Hesse. “Die Abhängigkeit der Viscosität von der Temperatur bie unterkühlten Flüssigkeiten”. In: *Zeitschrift für anorganische und allgemeine Chemie* 156.1 (1926), pp. 245–257. DOI: <https://doi.org/10.1002/zaac.19261560121>.
- [50] Gordon S. Fulcher. “ANALYSIS OF RECENT MEASUREMENTS OF THE VISCOSITY OF GLASSES”. In: *Journal of the American Ceramic Society* 8.6 (1925), pp. 339–355. DOI: <https://doi.org/10.1111/j.1151-2916.1925.tb16731.x>.
- [51] Giorgio Parisi, Pierfrancesco Urbani, and Francesco Zamponi. *Theory of Simple Glasses: Exact Solutions in Infinite Dimensions*. English. Hardcover. Cambridge University Press, Feb. 2020, p. 349. ISBN: 978-1107191075.
- [52] T. R. Kirkpatrick, D. Thirumalai, and P. G. Wolynes. “Scaling concepts for the dynamics of viscous liquids near an ideal glassy state”. In: *Phys. Rev. A* 40 (2 July 1989), pp. 1045–1054. DOI: 10.1103/PhysRevA.40.1045.
- [53] Marc Mézard and Giorgio Parisi. “A tentative replica study of the glass transition”. In: *Journal of Physics A: Mathematical and General* 29.20 (Oct. 1996), p. 6515. DOI: 10.1088/0305-4470/29/20/009.
- [54] Julian H Gibbs and Edmund A DiMarzio. “Nature of the glass transition and the glassy state”. In: *The Journal of Chemical Physics* 28.3 (1958), pp. 373–383.
- [55] Z Nussinov, N. B. Weingartner, and F. S. Nogueira. “The “glass transition” as a topological defect driven transition in a distribution of crystals and a prediction of a universal viscosity collapse”. In: *Topological Phase Transitions and New Developments*. Ed. by Lars Brink, Mike Gunn, et al. WORLD SCIENTIFIC, 2018. Chap. 6, pp. 61–79. DOI: 10.1142/11016.
- [56] Corey S. O’Hern, Leonardo E. Silbert, et al. “Jamming at zero temperature and zero applied stress: The epitome of disorder”. In: *Phys. Rev. E* 68 (1 July 2003), p. 011306. DOI: 10.1103/PhysRevE.68.011306.

- [57] Andrea J. Liu and Sidney R. Nagel. “The Jamming Transition and the Marginally Jammed Solid”. In: *Annual Review of Condensed Matter Physics* 1.1 (2010), pp. 347–369. DOI: 10.1146/annurev-conmatphys-070909-104045.
- [58] P. N. Pusey and W. van Megen. “Phase behaviour of concentrated suspensions of nearly hard colloidal spheres”. In: *Nature* 320.6060 (Mar. 1986), pp. 340–342. ISSN: 1476-4687. DOI: 10.1038/320340a0.
- [59] Eric R. Weeks. “Introduction to the Colloidal Glass Transition”. In: *ACS Macro Letters* 6.1 (Jan. 2017), pp. 27–34. DOI: 10.1021/acsmacrolett.6b00826.
- [60] Andrea Ninarello, Ludovic Berthier, and Daniele Coslovich. “Models and Algorithms for the Next Generation of Glass Transition Studies”. In: *Phys. Rev. X* 7 (2 June 2017), p. 021039. DOI: 10.1103/PhysRevX.7.021039.
- [61] Arthur K. Doolittle. “Studies in Newtonian Flow. II. The Dependence of the Viscosity of Liquids on Free-Space”. In: *Journal of Applied Physics* 22.12 (Dec. 1951), pp. 1471–1475. ISSN: 0021-8979. DOI: 10.1063/1.1699894.
- [62] Zhengdong Cheng, Jixiang Zhu, et al. “Nature of the divergence in low shear viscosity of colloidal hard-sphere dispersions”. In: *Phys. Rev. E* 65 (4 Apr. 2002), p. 041405. DOI: 10.1103/PhysRevE.65.041405.
- [63] Boris D. Lubachevski and Frank H. Stillinger. “Geometric Properties of Random Disk Packings”. In: *Journal of Statistical Physics* 60 (5/6 Mar. 1990), pp. 561–583. DOI: 10.1007/BF01025983.
- [64] G D Scott and D M Kilgour. “The density of random close packing of spheres”. In: *Journal of Physics D: Applied Physics* 2.6 (June 1969), pp. 863–866. DOI: 10.1088/0022-3727/2/6/311.
- [65] Ning Xu, Jerzy Blawdziewicz, and Corey S. O’Hern. “Random close packing revisited: Ways to pack frictionless disks”. In: *Phys. Rev. E* 71 (6 June 2005), p. 061306. DOI: 10.1103/PhysRevE.71.061306.
- [66] Atsushi Ikeda, Ludovic Berthier, and Peter Sollich. “Unified study of glass and jamming rheology in soft particle systems”. In: *Phys. Rev. Lett.* 109 (1 July 2012), p. 018301. DOI: 10.1103/PhysRevLett.109.018301.
- [67] Keishi Gotoh and John L. Finney. “Statistical geometrical approach to random packing density of equal spheres”. In: *Nature* 252.5480 (Nov. 1974), pp. 202–205. ISSN: 1476-4687. DOI: 10.1038/252202a0.
- [68] James G. Berryman. “Random close packing of hard spheres and disks”. In: *Phys. Rev. A* 27 (2 Feb. 1983), pp. 1053–1061. DOI: 10.1103/PhysRevA.27.1053.
- [69] O. Pouliquen, M. Nicolas, and P. D. Weidman. “Crystallization of non-Brownian Spheres under Horizontal Shaking”. In: *Phys. Rev. Lett.* 79 (19 Nov. 1997), pp. 3640–3643. DOI: 10.1103/PhysRevLett.79.3640.
- [70] S. Torquato, T. M. Truskett, and P. G. Debenedetti. “Is Random Close Packing of Spheres Well Defined?” In: *Phys. Rev. Lett.* 84 (10 Mar. 2000), pp. 2064–2067. DOI: 10.1103/PhysRevLett.84.2064.

- [71] Andrea J. Liu and Sidney R. Nagel. “Jamming is not just cool any more”. In: *Nature* 396.6706 (Nov. 1998), pp. 21–22. ISSN: 1476-4687. DOI: 10.1038/23819.
- [72] M. E. J. Newman and Cristopher Moore. “Glassy dynamics and aging in an exactly solvable spin model”. In: *Phys. Rev. E* 60 (5 Oct. 1999), pp. 5068–5072. DOI: 10.1103/PhysRevE.60.5068.
- [73] Sergey Bravyi, Bernhard Leemhuis, and Barbara M. Terhal. “Topological order in an exactly solvable 3D spin model”. In: *Annals of Physics* 326.4 (2011), pp. 839–866. ISSN: 0003-4916. DOI: <https://doi.org/10.1016/j.aop.2010.11.002>.
- [74] Sergey Bravyi and Jeongwan Haah. “Quantum Self-Correction in the 3D Cubic Code Model”. In: *Phys. Rev. Lett.* 111 (20 Oct. 2013), p. 200501. DOI: 10.1103/PhysRevLett.111.200501.
- [75] Claudio Castelnovo and Claudio Chamon. “Topological quantum glassiness”. In: *Philosophical Magazine* 92.1-3 (2012), pp. 304–323. DOI: 10.1080/14786435.2011.609152.
- [76] Beni Yoshida. “Exotic topological order in fractal spin liquids”. In: *Phys. Rev. B* 88 (12 Sept. 2013), p. 125122. DOI: 10.1103/PhysRevB.88.125122.
- [77] Sagar Vijay, Jeongwan Haah, and Liang Fu. “A new kind of topological quantum order: A dimensional hierarchy of quasiparticles built from stationary excitations”. In: *Phys. Rev. B* 92 (23 Dec. 2015), p. 235136. DOI: 10.1103/PhysRevB.92.235136.
- [78] Sagar Vijay, Jeongwan Haah, and Liang Fu. “Fracton topological order, generalized lattice gauge theory, and duality”. In: *Phys. Rev. B* 94 (23 Dec. 2016), p. 235157. DOI: 10.1103/PhysRevB.94.235157.
- [79] Jeongwan Haah. “Local stabilizer codes in three dimensions without string logical operators”. In: *Phys. Rev. A* 83 (4 Apr. 2011), p. 042330. DOI: 10.1103/PhysRevA.83.042330.
- [80] Michael Pretko, Xie Chen, and Yizhi You. “Fracton phases of matter”. In: *International Journal of Modern Physics A* 35.06 (2020), p. 2030003. DOI: 10.1142/S0217751X20300033.
- [81] “Fractons”. In: *Annual Review of Condensed Matter Physics* 10.1 (2019), pp. 295–313. DOI: 10.1146/annurev-conmatphys-031218-013604.
- [82] Cenke Xu. “Gapless bosonic excitation without symmetry breaking: An algebraic spin liquid with soft gravitons”. In: *Phys. Rev. B* 74 (22 Dec. 2006), p. 224433. DOI: 10.1103/PhysRevB.74.224433.
- [83] Cenke Xu. “Emergent gravity at a Lifshitz point from a Bose liquid on the lattice”. In: *Phys. Rev. D* 81 (10 May 2010), p. 104033. DOI: 10.1103/PhysRevD.81.104033.
- [84] Michael Pretko. “Subdimensional particle structure of higher rank $U(1)$ spin liquids”. In: *Phys. Rev. B* 95 (11 Mar. 2017), p. 115139. DOI: 10.1103/PhysRevB.95.115139.

- [85] Michael Pretko. “Generalized electromagnetism of subdimensional particles: A spin liquid story”. In: *Phys. Rev. B* 96 (3 July 2017), p. 035119. DOI: 10.1103/PhysRevB.96.035119.
- [86] Michael Pretko. “Emergent gravity of fractons: Mach’s principle revisited”. In: *Phys. Rev. D* 96 (2 July 2017), p. 024051. DOI: 10.1103/PhysRevD.96.024051.
- [87] Hao Song, Janik Schönmeier-Kromer, et al. “Optimal Thresholds for Fracton Codes and Random Spin Models with Subsystem Symmetry”. In: *Phys. Rev. Lett.* 129 (23 Oct. 2022), p. 230502. DOI: 10.1103/PhysRevLett.129.230502.
- [88] A. T. Schmitz, Han Ma, et al. “Recoverable information and emergent conservation laws in fracton stabilizer codes”. In: *Phys. Rev. B* 97 (13 Apr. 2018), p. 134426. DOI: 10.1103/PhysRevB.97.134426.
- [89] Alex Rasmussen, Yi-Zhuang You, and Cenke Xu. *Stable Gapless Bose Liquid Phases without any Symmetry*. May 2016. DOI: 10.1103/PhysRevD.81.104033.
- [90] Han Ma, Michael Hermele, and Xie Chen. “Fracton topological order from the Higgs and partial-confinement mechanisms of rank-two gauge theory”. In: *Phys. Rev. B* 98 (3 July 2018), p. 035111. DOI: 10.1103/PhysRevB.98.035111.
- [91] Daniel Bulmash and Maissam Barkeshli. “Higgs mechanism in higher-rank symmetric $U(1)$ gauge theories”. In: *Phys. Rev. B* 97 (23 June 2018), p. 235112. DOI: 10.1103/PhysRevB.97.235112.
- [92] Han Yan, Owen Benton, et al. “Rank-2 $U(1)$ Spin Liquid on the Breathing Pyrochlore Lattice”. In: *Phys. Rev. Lett.* 124 (12 Mar. 2020), p. 127203. DOI: 10.1103/PhysRevLett.124.127203.
- [93] Owen Benton and Roderich Moessner. “Topological Route to New and Unusual Coulomb Spin Liquids”. In: *Phys. Rev. Lett.* 127 (10 Sept. 2021), p. 107202. DOI: 10.1103/PhysRevLett.127.107202.
- [94] Grace M. Sommers, Benedikt Placke, et al. “From hard spheres to hard-core spins”. In: *Phys. Rev. B* 103 (10 Mar. 2021), p. 104407. DOI: 10.1103/PhysRevB.103.104407.
- [95] Benedikt Placke, Grace M. Sommers, et al. “Arresting dynamics in hardcore spin models”. In: *Phys. Rev. B* 107 (18 May 2023), p. L180302. DOI: 10.1103/PhysRevB.107.L180302.
- [96] Nigel Goldenfeld. *Lectures on phase transitions and the Renormalization Group*. Vol. 81. Addison-Wesley, May 1992, p. 104033. DOI: 10.1103/PhysRevD.81.104033.
- [97] Piotr Miłoś and Ron Peled. “Delocalization of two-dimensional random surfaces with hard-core constraints”. In: *Communications in Mathematical Physics* 340.1 (10 Apr. 2015), pp. 1–46. ISSN: 14320916. DOI: 10.1007/s00220-015-2419-4.
- [98] Ron Peled. *Emergent gravity at a Lifshitz point from a Bose liquid on the lattice*. (private communication). May 2010. DOI: 10.1103/PhysRevD.81.104033.

- [99] Wolfgang Bietenholz, Urs Gerber, and Fernando G. Rejón-Barrera. “Berezinskii-Kosterlitz-Thouless transition with a constraint lattice action”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2013.12 (10 May 2013), P12009. ISSN: 17425468. DOI: 10.1088/1742-5468/2013/12/P12009.
- [100] Steven Atkinson, Frank H. Stillinger, and Salvatore Torquato. “Existence of isostatic, maximally random jammed monodisperse hard-disk packings”. In: *Proceedings of the National Academy of Sciences* 111.52 (10 Dec. 2014), pp. 18436–18441. ISSN: 0027-8424. DOI: 10.1073/pnas.1408371112.
- [101] B. Yurke, A. N. Pargellis, et al. “Coarsening dynamics of the XY model”. In: *Phys. Rev. E* 47 (3 Mar. 1993), pp. 1525–1530. DOI: 10.1103/PhysRevE.47.1525.
- [102] Lars Melwyn Jensen, Beom Jun Kim, and Petter Minnhagen. “Dynamic critical exponent of two-, three-, and four-dimensional XY models with relaxational and resistively shunted junction dynamics”. In: *Phys. Rev. B* 61 (22 June 2000), pp. 15412–15428. DOI: 10.1103/PhysRevB.61.15412.
- [103] Benedikt Placke, Owen Benton, and Roderich Moessner. *Ising Fracton Spin Liquid on the Honeycomb Lattice*. arXiv:2306.13151. May 2023. DOI: 10.48550/arXiv.2306.13151.
- [104] Michael Pretko and Leo Radzihovsky. “Fracton-Elasticity Duality”. In: *Phys. Rev. Lett.* 120 (19 May 2018), p. 195301. DOI: 10.1103/PhysRevLett.120.195301.
- [105] John Sous and Michael Pretko. “Fractons from frustration in hole-doped antiferromagnets”. In: *npj Quantum Materials* 5.1 (10 Nov. 2020), p. 81. DOI: 10.1038/s41535-020-00278-2.
- [106] Yizhi You, Zhen Bi, and Michael Pretko. “Emergent fractons and algebraic quantum liquid from plaquette melting transitions”. In: *Phys. Rev. Res.* 2 (1 Feb. 2020), p. 013162. DOI: 10.1103/PhysRevResearch.2.013162.
- [107] Michael Pretko. “Electric circuit realizations of fracton physics”. In: *Phys. Rev. B* 100 (24 Dec. 2019), p. 245103. DOI: 10.1103/PhysRevB.100.245103.
- [108] Max Hering, Han Yan, and Johannes Reuther. “Fracton excitations in classical frustrated kagome spin models”. In: *Phys. Rev. B* 104 (6 Aug. 2021), p. 064406. DOI: 10.1103/PhysRevB.104.064406.
- [109] Han Yan and Johannes Reuther. “Low-energy structure of spiral spin liquids”. In: *Phys. Rev. Res.* 4 (2 June 2022), p. 023175. DOI: 10.1103/PhysRevResearch.4.023175.
- [110] Nayan E. Myerson-Jain, Stephen Yan, et al. “Construction of Fractal Order and Phase Transition with Rydberg Atoms”. In: *Phys. Rev. Lett.* 128 (1 Jan. 2022), p. 017601. DOI: 10.1103/PhysRevLett.128.017601.
- [111] Krzysztof Giergiel, Ruben Lier, et al. “Bose-Hubbard realization of fracton defects”. In: *Phys. Rev. Res.* 4 (2 May 2022), p. 023151. DOI: 10.1103/PhysRevResearch.4.023151.
- [112] SangEun Han, Adarsh S. Patri, and Yong Baek Kim. “Realization of fractonic quantum phases in the breathing pyrochlore lattice”. In: *Phys. Rev. B* 105 (23 June 2022), p. 235120. DOI: 10.1103/PhysRevB.105.235120.

-
- [113] Han Yan, Owen Benton, et al. *Classification of Classical Spin Liquids: Typology and Resulting Landscape*. May 2023. DOI: 10.1103/PhysRevD.81.104033.
- [114] Han Yan, Owen Benton, et al. *Classification of Classical Spin Liquids: Detailed Formalism and Suite of Examples*. May 2023. DOI: 10.1103/PhysRevD.81.104033.
- [115] Abhinav Prem, Sagar Vijay, et al. “Pinch point singularities of tensor spin liquids”. In: *Phys. Rev. B* 98 (16 Oct. 2018), p. 165140. DOI: 10.1103/PhysRevB.98.165140.
- [116] S. V. Isakov, K. Gregor, et al. “Dipolar Spin Correlations in Classical Pyrochlore Magnets”. In: *Phys. Rev. Lett.* 93.16, 167204 (10 Oct. 2004), p. 167204. DOI: 10.1103/PhysRevLett.93.167204.
- [117] Itiro Syôzi. “Statistics of Kagomé Lattice”. In: *Progress of Theoretical Physics* 6.3 (10 June 1951), pp. 306–308. ISSN: 0033-068X. DOI: 10.1143/ptp/6.3.306.
- [118] K. Kanô and S. Naya. “Antiferromagnetism. The Kagomé Ising Net”. In: *Progress of Theoretical Physics* 10.2 (10 Aug. 1953), pp. 158–172. DOI: 10.1143/ptp/10.2.158.
- [119] D. A. Garanin and Benjamin Canals. “Classical spin liquid: Exact solution for the infinite-component antiferromagnetic model on the kagomé lattice”. In: *Phys. Rev. B* 59.1 (10 Jan. 1999), pp. 443–456. DOI: 10.1103/PhysRevB.59.443.
- [120] C. Castelnovo, R. Moessner, and S. L. Sondhi. “Magnetic monopoles in spin ice”. In: *Nature* 451.7174 (10 Jan. 2008), pp. 42–45. ISSN: 1476-4687. DOI: 10.1038/nature06433.
- [121] R. Youngblood, J. D. Axe, and B. M. McCoy. “Correlations in ice-rule ferroelectrics”. In: *Phys. Rev. B* 21 (11 June 1980), pp. 5212–5220. DOI: 10.1103/PhysRevB.21.5212.
- [122] C. L. Henley. “Power-law spin correlations in pyrochlore antiferromagnets”. In: *Phys. Rev. B* 71 (1 Jan. 2005), p. 014424. DOI: 10.1103/PhysRevB.71.014424.
- [123] Werner Krauth. “Introduction To Monte Carlo Algorithms”. In: *Phys. Rev. D* 81 (10 May 1996), p. 104033. DOI: 10.48550/arXiv.cond-mat/9612186.
- [124] Helmut G Katzgraber, Simon Trebst, et al. “Feedback-optimized parallel tempering Monte Carlo”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2006.03 (10 Mar. 2006), P03018. DOI: 10.1088/1742-5468/2006/03/P03018.
- [125] Ludovic Berthier and David R. Reichman. “Modern computational studies of the glass transition”. In: *Nature Reviews Physics* 5.2 (10 Feb. 2023), pp. 102–116. ISSN: 2522-5820. DOI: 10.1038/s42254-022-00548-x.
- [126] Nils Niggemann, Yasir Iqbal, and Johannes Reuther. “Quantum Effects on Unconventional Pinch Point Singularities”. In: *Phys. Rev. Lett.* 130 (19 May 2023), p. 196601. DOI: 10.1103/PhysRevLett.130.196601.

- [127] David Deutsch and Richard Jozsa. “Rapid solution of problems by quantum computation”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439.1907 (10 May 1992), pp. 553–558. DOI: 10.1098/rspa.1992.0167.
- [128] P.W. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. Vol. 81. American Physical Society, May 1994, pp. 124–134. DOI: 10.1109/SFCS.1994.365700.
- [129] Sergey Bravyi, David Gosset, and Robert König. “Quantum advantage with shallow circuits”. In: *Science* 362.6412 (10 May 2018), pp. 308–311. DOI: 10.1126/science.aar3106.
- [130] Lov K. Grover. “A Fast Quantum Mechanical Algorithm for Database Search”. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. Vol. 81. STOC '96. New York, NY, USA: Association for Computing Machinery, May 1996, 212–219. ISBN: 0897917855. DOI: 10.1145/237814.237866.
- [131] Ethan Bernstein and Umesh Vazirani. “Quantum Complexity Theory”. In: *SIAM Journal on Computing* 26.5 (10 May 1997), pp. 1411–1473. DOI: 10.1137/S0097539796300921.
- [132] Sergio Boixo, Sergei V. Isakov, et al. “Characterizing quantum supremacy in near-term devices”. In: *Nature Physics* 14.6 (10 June 2018), pp. 595–600. ISSN: 1745-2481. DOI: 10.1038/s41567-018-0124-x.
- [133] Arnold Schönhage. “On the Power of Random Access Machines”. In: *International Colloquium on Automata, Languages and Programming*. Vol. 81. American Physical Society, May 1979, p. 104033. DOI: 10.1007/3-540-09510-1_42.
- [134] Peter W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52 (4 Oct. 1995), R2493–R2496. DOI: 10.1103/PhysRevA.52.R2493.
- [135] R. Arratia and L. Gordon. “Tutorial on large deviations for the binomial distribution”. In: *Bulletin of Mathematical Biology* 51.1 (10 Jan. 1989), pp. 125–131. ISSN: 1522-9602. DOI: 10.1007/BF02458840.
- [136] Eric Dennis, Alexei Kitaev, et al. “Topological quantum memory”. In: *Journal of Mathematical Physics* 43.9 (10 Aug. 2002), pp. 4452–4505. ISSN: 0022-2488. DOI: 10.1063/1.1499754.
- [137] Barbara M. Terhal. “Quantum error correction for quantum memories”. In: *Rev. Mod. Phys.* 87 (2 Apr. 2015), pp. 307–346. DOI: 10.1103/RevModPhys.87.307.
- [138] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. 10th. Vol. 81. USA: Cambridge University Press, May 2011, p. 104033. ISBN: 1107002176. DOI: 10.1103/PhysRevD.81.104033.
- [139] Jack Edmonds. “Paths, Trees, and Flowers”. In: *Canadian Journal of Mathematics* 17 (10 May 1965), 449–467. DOI: 10.4153/CJM-1965-045-4.

- [140] Austin G. Fowler. “Minimum Weight Perfect Matching of Fault-Tolerant Topological Quantum Error Correction in Average $O(1)$ Parallel Time”. In: *Quantum Info. Comput.* 15.1–2 (10 Jan. 2015), 145–158. ISSN: 1533-7146. DOI: 10.1103/PhysRevD.81.104033.
- [141] Hidetoshi Nishimori. “Internal Energy, Specific Heat and Correlation Function of the Bond-Random Ising Model”. In: *Progress of Theoretical Physics* 66.4 (10 May 1981), pp. 1169–1181. DOI: 10.1143/PTP.66.1169.
- [142] P.W. Kasteleyn. “The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice”. In: *Physica* 27.12 (10 May 1961), pp. 1209–1225. ISSN: 0031-8914. DOI: [https://doi.org/10.1016/0031-8914\(61\)90063-5](https://doi.org/10.1016/0031-8914(61)90063-5).
- [143] H. N. V. Temperley and Michael E. Fisher. “Dimer problem in statistical mechanics-an exact result”. In: *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics* 6.68 (10 May 1961), pp. 1061–1063. DOI: 10.1080/14786436108243366.
- [144] P. W. Kasteleyn. “Dimer Statistics and Phase Transitions”. In: *Journal of Mathematical Physics* 4.2 (10 Dec. 1963), pp. 287–293. ISSN: 0022-2488. DOI: 10.1063/1.1703953.
- [145] Thomas Vojta. “Rare region effects at classical, quantum and nonequilibrium phase transitions”. In: *Journal of Physics A: Mathematical and General* 39.22 (10 May 2006), R143. DOI: 10.1088/0305-4470/39/22/R01.
- [146] Amnon Aharony and A. Brooks Harris. “Absence of Self-Averaging and Universal Fluctuations in Random Systems near Critical Points”. In: *Phys. Rev. Lett.* 77 (18 Oct. 1996), pp. 3700–3703. DOI: 10.1103/PhysRevLett.77.3700.
- [147] Sora Cho and Matthew P. A. Fisher. “Criticality in the two-dimensional random-bond Ising model”. In: *Phys. Rev. B* 55 (2 Jan. 1997), pp. 1025–1031. DOI: 10.1103/PhysRevB.55.1025.
- [148] F. Merz and J. T. Chalker. “Two-dimensional random-bond Ising model, free fermions, and the network model”. In: *Phys. Rev. B* 65 (5 Jan. 2002), p. 054425. DOI: 10.1103/PhysRevB.65.054425.
- [149] Chenyang Wang, Jim Harrington, and John Preskill. “Confinement-Higgs transition in a disordered gauge theory and the accuracy threshold for quantum memory”. In: *Annals of Physics* 303.1 (10 May 2003), pp. 31–58. ISSN: 0003-4916. DOI: [https://doi.org/10.1016/S0003-4916\(02\)00019-2](https://doi.org/10.1016/S0003-4916(02)00019-2).
- [150] Helmut G. Katzgraber, H. Bombin, and M. A. Martin-Delgado. “Error Threshold for Color Codes and Random Three-Body Ising Models”. In: *Phys. Rev. Lett.* 103 (9 Aug. 2009), p. 090501. DOI: 10.1103/PhysRevLett.103.090501.
- [151] Aleksander Kubica, Michael E. Beverland, et al. “Three-Dimensional Color Code Thresholds via Statistical-Mechanical Mapping”. In: *Phys. Rev. Lett.* 120 (18 May 2018), p. 180501. DOI: 10.1103/PhysRevLett.120.180501.
- [152] Hao Song, Janik Schönmeier-Kromer, et al. “Optimal Thresholds for Fracton Codes and Random Spin Models with Subsystem Symmetry”. In: *Phys. Rev. Lett.* 129 (23 Nov. 2022), p. 230502. DOI: 10.1103/PhysRevLett.129.230502.

- [153] H. Bombin, Ruben S. Andrist, et al. “Strong Resilience of Topological Codes to Depolarization”. In: *Phys. Rev. X* 2 (2 Apr. 2012), p. 021004. DOI: 10.1103/PhysRevX.2.021004.
- [154] Christopher T. Chubb and Steven T. Flammia. “Statistical mechanical models for quantum codes with correlated noise”. In: *Phys. Rev. D* 81 (10 May 2010), p. 104033. DOI: 10.1103/PhysRevD.81.104033.
- [155] Florian Venn, Jan Behrends, and Benjamin Béri. *Coherent error threshold for surface codes from Majorana delocalization*. May 2022. DOI: 10.1103/PhysRevD.81.104033.
- [156] H. Bombin and M. A. Martin-Delgado. “Topological Quantum Distillation”. In: *Phys. Rev. Lett.* 97 (18 Oct. 2006), p. 180501. DOI: 10.1103/PhysRevLett.97.180501.
- [157] Héctor Bombín. “Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes”. In: *New Journal of Physics* 17.8 (10 Aug. 2015), p. 083002. DOI: 10.1088/1367-2630/17/8/083002.
- [158] Aleksander Kubica and Michael E. Beverland. “Universal transversal gates with color codes: A simplified approach”. In: *Phys. Rev. A* 91 (3 Mar. 2015), p. 032330. DOI: 10.1103/PhysRevA.91.032330.
- [159] Nikolas P. Breuckmann, Benedikt Placke, and Ananda Roy. “Critical properties of the Ising model in hyperbolic space”. In: *Phys. Rev. E* 101 (2 Feb. 2020), p. 022124. DOI: 10.1103/PhysRevE.101.022124.
- [160] Benedikt Placke and Nikolas P. Breuckmann. “Random-bond Ising model and its dual in hyperbolic spaces”. In: *Phys. Rev. E* 107 (2 Feb. 2023), p. 024125. DOI: 10.1103/PhysRevE.107.024125.
- [161] R Rietman, B Nienhuis, and J Oitmaa. “The Ising model on hyperlattices”. In: *Journal of Physics A: Mathematical and General* 25.24 (10 Dec. 1992), pp. 6577–6592. DOI: 10.1088/0305-4470/25/24/012.
- [162] Nikolas P Breuckmann. “Homological quantum codes beyond the toric code”. PhD thesis. RWTH Aachen University, May 2017, p. 104033. DOI: 10.1103/PhysRevD.81.104033.
- [163] Nikolas P Breuckmann and Barbara M Terhal. “Constructions and noise threshold of hyperbolic surface codes”. In: *IEEE transactions on Information Theory* 62.6 (10 May 2016), pp. 3731–3744. DOI: 10.1103/PhysRevD.81.104033.
- [164] Martin Macaj, Jozef Siran, and Mária Ipolyiová. “Injectivity radius of representations of triangle groups and planar width of regular hypermaps.” In: *Ars Math. Contemp.* 1.2 (10 May 2008), pp. 223–241. DOI: 10.1103/PhysRevD.81.104033.
- [165] Judith Flagg Moran. “The growth rate and balance of homogeneous tilings in the hyperbolic plane”. In: *Discrete Mathematics* 173.1-3 (10 May 1997), pp. 151–186. DOI: 10.1103/PhysRevD.81.104033.
- [166] Nikolas P Breuckmann, Christophe Vuillot, et al. “Hyperbolic and semi-hyperbolic surface codes for quantum storage”. In: *Quantum Science and Technology* 2.3 (10 May 2017), p. 035007. DOI: 10.1103/PhysRevD.81.104033.

-
- [167] Curtis G. Callan and Frank Wilczek. “Infrared behavior at negative curvature”. In: *Nuclear Physics B* 340.2 (10 May 1990), pp. 366–386. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(90\)90451-I](https://doi.org/10.1016/0550-3213(90)90451-I).
- [168] Alicia J. Kollár, Mattias Fitzpatrick, and Andrew A. Houck. “Hyperbolic lattices in circuit quantum electrodynamics”. In: *Nature* 571.7763 (10 July 2019), pp. 45–50. ISSN: 1476-4687. DOI: [10.1038/s41586-019-1348-3](https://doi.org/10.1038/s41586-019-1348-3).
- [169] Przemyslaw Bienias, Igor Boettcher, et al. “Circuit Quantum Electrodynamics in Hyperbolic Space: From Photon Bound States to Frustrated Spin Models”. In: *Phys. Rev. Lett.* 128 (1 Jan. 2022), p. 013601. DOI: [10.1103/PhysRevLett.128.013601](https://doi.org/10.1103/PhysRevLett.128.013601).
- [170] Joseph Maciejko and Steven Rayan. “Hyperbolic band theory”. In: *Science Advances* 7.36 (10 May 2021), eabe9170. DOI: [10.1126/sciadv.abe9170](https://doi.org/10.1126/sciadv.abe9170).
- [171] Igor Boettcher, Alexey V. Gorshkov, et al. “Crystallography of hyperbolic lattices”. In: *Phys. Rev. B* 105 (12 Mar. 2022), p. 125118. DOI: [10.1103/PhysRevB.105.125118](https://doi.org/10.1103/PhysRevB.105.125118).
- [172] Adil Attar and Igor Boettcher. “Selberg trace formula in hyperbolic band theory”. In: *Phys. Rev. E* 106 (3 Sept. 2022), p. 034114. DOI: [10.1103/PhysRevE.106.034114](https://doi.org/10.1103/PhysRevE.106.034114).
- [173] Kazuki Ikeda, Shoto Aoki, and Yoshiyuki Matsuki. “Hyperbolic band theory under magnetic field and Dirac cones on a higher genus surface”. In: *Journal of Physics: Condensed Matter* 33.48 (10 Sept. 2021), p. 485602. DOI: [10.1088/1361-648X/ac24c4](https://doi.org/10.1088/1361-648X/ac24c4).
- [174] Patrick M. Lenggenhager, Alexander Stegmaier, et al. “Simulating hyperbolic space on a circuit board”. In: *Nature Communications* 13.1 (10 July 2022), p. 4373. ISSN: 2041-1723. DOI: [10.1038/s41467-022-32042-4](https://doi.org/10.1038/s41467-022-32042-4).
- [175] R Krcmar, A Gendiar, et al. “Ising model on a hyperbolic lattice studied by the corner transfer matrix renormalization group method”. In: *Journal of Physics A: Mathematical and Theoretical* 41.12 (10 Mar. 2008), p. 125001. DOI: [10.1088/1751-8113/41/12/125001](https://doi.org/10.1088/1751-8113/41/12/125001).
- [176] Karim Mnasri, Bhilahari Jeevanesan, and Jörg Schmalian. “Critical phenomena in hyperbolic space”. In: *Phys. Rev. B* 92 (13 Oct. 2015), p. 134423. DOI: [10.1103/PhysRevB.92.134423](https://doi.org/10.1103/PhysRevB.92.134423).
- [177] Yi Jiang, Ilya Dumer, et al. “Duality and free energy analyticity bounds for few-body Ising models with extensive homology rank”. In: *Journal of Mathematical Physics* 60.8 (10 May 2019), p. 083302. DOI: [10.1103/PhysRevD.81.104033](https://doi.org/10.1103/PhysRevD.81.104033).
- [178] C. Chris Wu. “Ising models on hyperbolic graphs”. In: *Journal of Statistical Physics* 85.1 (10 Oct. 1996), pp. 251–259. ISSN: 1572-9613. DOI: [10.1007/BF02175564](https://doi.org/10.1007/BF02175564).
- [179] C. Chris Wu. “Ising Models on Hyperbolic Graphs II”. In: *Journal of Statistical Physics* 100.5 (10 Sept. 2000), pp. 893–904. ISSN: 1572-9613. DOI: [10.1023/A:1018763008810](https://doi.org/10.1023/A:1018763008810).
- [180] Daniel S Freed and Constantin Teleman. “Topological dualities in the Ising model”. In: *arXiv preprint arXiv:1806.00008* 81 (10 May 2018), p. 104033. DOI: [10.1103/PhysRevD.81.104033](https://doi.org/10.1103/PhysRevD.81.104033).

- [181] A.M. Polyakov. *Gauge Fields and Strings*. 1st ed. Vol. 81. Routledge, May 1987, p. 104033. DOI: 10.1201/9780203755082.
- [182] VI. S. Dotsenko and A. M. Polyakov. “Fermion Representations for the 2D and 3D Ising Models”. In: *Conformal Field Theory and Solvable Lattice Models* 81 (10 May 1988), 171–203. ISSN: 0920-1971. DOI: 10.2969/aspm/01610171.
- [183] H. A. Kramers and G. H. Wannier. “Statistics of the Two-Dimensional Ferromagnet. Part I”. In: *Phys. Rev.* 60 (3 Aug. 1941), pp. 252–262. DOI: 10.1103/PhysRev.60.252.
- [184] J. Oitmaa, C. Hamer, et al. *Series Expansion Methods for Strongly Interacting Lattice Models*. Vol. 81. Cambridge University Press, May 2006, p. 104033. ISBN: 9780521842426. DOI: 10.1103/PhysRevD.81.104033.
- [185] Koujin Takeda, Tomohiro Sasamoto, and Hidetoshi Nishimori. “Exact location of the multicritical point for finite-dimensional spin glasses: a conjecture”. In: *Journal of Physics A: Mathematical and General* 38.17 (10 May 2005), p. 3751. DOI: 10.1103/PhysRevD.81.104033.
- [186] J. M. Carlson, J. T. Chayes, et al. “Bethe lattice spin glass: The effects of a ferromagnetic bias and external fields. II. Magnetized spin-glass phase and the de Almeida-Thouless line”. In: *Journal of Statistical Physics* 61.5 (10 Dec. 1990), pp. 1069–1084. ISSN: 1572-9613. DOI: 10.1007/BF01014365.
- [187] D. J. Thouless. “Spin-Glass on a Bethe Lattice”. In: *Phys. Rev. Lett.* 56 (10 Mar. 1986), pp. 1082–1085. DOI: 10.1103/PhysRevLett.56.1082.
- [188] R. Botet, R. Jullien, and P. Pfeuty. “Size Scaling for Infinitely Coordinated Systems”. In: *Phys. Rev. Lett.* 49 (7 Aug. 1982), pp. 478–481. DOI: 10.1103/PhysRevLett.49.478.
- [189] Hiroyuki Shima and Yasunori Sakaniwa. “Geometric effects on critical behaviours of the Ising model”. In: *Journal of Physics A: Mathematical and General* 39.18 (10 Apr. 2006), pp. 4921–4933. DOI: 10.1088/0305-4470/39/18/010.
- [190] Fugao Wang and D. P. Landau. “Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States”. In: *Phys. Rev. Lett.* 86 (10 Mar. 2001), pp. 2050–2053. DOI: 10.1103/PhysRevLett.86.2050.
- [191] Fugao Wang and D. P. Landau. “Determining the density of states for classical statistical models: A random walk algorithm to produce a flat histogram”. In: *Phys. Rev. E* 64 (5 Oct. 2001), p. 056101. DOI: 10.1103/PhysRevE.64.056101.
- [192] B. J. Schulz, K. Binder, et al. “Avoiding boundary effects in Wang-Landau sampling”. In: *Phys. Rev. E* 67 (6 June 2003), p. 067102. DOI: 10.1103/PhysRevE.67.067102.
- [193] R. E. Belardinelli and V. D. Pereyra. “Wang-Landau algorithm: A theoretical analysis of the saturation of the error”. In: *The Journal of Chemical Physics* 127.18 (10 May 2007), p. 184105. DOI: 10.1063/1.2803061.
- [194] Daniel S. Rokhsar and Steven A. Kivelson. “Superconductivity and the Quantum Hard-Core Dimer Gas”. In: *Phys. Rev. Lett.* 61 (20 Nov. 1988), pp. 2376–2379. DOI: 10.1103/PhysRevLett.61.2376.

- [195] R. Moessner and S. L. Sondhi. “Resonating Valence Bond Phase in the Triangular Lattice Quantum Dimer Model”. In: *Phys. Rev. Lett.* 86 (9 Feb. 2001), pp. 1881–1884. DOI: 10.1103/PhysRevLett.86.1881.
- [196] 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”. In: *Phys. Rev. B* 69 (6 Feb. 2004), p. 064404. DOI: 10.1103/PhysRevB.69.064404.
- [197] C. Broholm, R. J. Cava, et al. “Quantum spin liquids”. In: *Science* 367.6475 (10 May 2020), eaay0668. DOI: 10.1126/science.aay0668.
- [198] P. W. Anderson. “Ordering and Antiferromagnetism in Ferrites”. In: *Phys. Rev.* 102 (4 May 1956), pp. 1008–1013. DOI: 10.1103/PhysRev.102.1008.
- [199] J. D. Bernal and R. H. Fowler. “A Theory of Water and Ionic Solution, with Particular Reference to Hydrogen and Hydroxyl Ions”. In: *The Journal of Chemical Physics* 1.8 (10 Nov. 2004), pp. 515–548. ISSN: 0021-9606. DOI: 10.1063/1.1749327.
- [200] Linus Pauling. “The Structure and Entropy of Ice and of Other Crystals with Some Randomness of Atomic Arrangement”. In: *Journal of the American Chemical Society* 57.12 (10 Dec. 1935), pp. 2680–2684. ISSN: 0002-7863. DOI: 10.1021/ja01315a102.
- [201] J. F. Nagle. “Lattice Statistics of Hydrogen Bonded Crystals. I. The Residual Entropy of Ice”. In: *Journal of Mathematical Physics* 7.8 (10 Dec. 1966), pp. 1484–1491. ISSN: 0022-2488. DOI: 10.1063/1.1705058.
- [202] R. R. P. Singh and J. Oitmaa. “Corrections to Pauling residual entropy and single tetrahedron based approximations for the pyrochlore lattice Ising antiferromagnet”. In: *Phys. Rev. B* 85 (14 Apr. 2012), p. 144414. DOI: 10.1103/PhysRevB.85.144414.
- [203] Steven T. Bramwell and Michel J. P. Gingras. “Spin Ice State in Frustrated Magnetic Pyrochlore Materials”. In: *Science* 294.5546 (10 May 2001), pp. 1495–1501. DOI: 10.1126/science.1064761.
- [204] M. J. Harris, S. T. Bramwell, et al. “Geometrical Frustration in the Ferromagnetic Pyrochlore $\text{Ho}_2\text{Ti}_2\text{O}_7$ ”. In: *Phys. Rev. Lett.* 79 (13 Sept. 1997), pp. 2554–2557. DOI: 10.1103/PhysRevLett.79.2554.
- [205] A. P. Ramirez, A. Hayashi, et al. “Zero-point entropy in ‘spin ice’”. In: *Nature* 399.6734 (10 May 1999), pp. 333–335. ISSN: 1476-4687. DOI: 10.1038/20619.
- [206] T. Fennell, P. P. Deen, et al. “Magnetic Coulomb Phase in the Spin Ice $\text{Ho}_2\text{Ti}_2\text{O}_7$ ”. In: *Science* 326.5951 (10 May 2009), pp. 415–417. DOI: 10.1126/science.1177582.
- [207] D. J. P. Morris, D. A. Tennant, et al. “Dirac Strings and Magnetic Monopoles in the Spin Ice $\text{Dy}_2\text{Ti}_2\text{O}_7$ ”. In: *Science* 326.5951 (10 May 2009), pp. 411–414. DOI: 10.1126/science.1178868.
- [208] Jonathan N. Hallén, Santiago A. Grigera, et al. “Dynamical fractal and anomalous noise in a clean magnetic crystal”. In: *Science* 378.6625 (10 May 2022), pp. 1218–1221. DOI: 10.1126/science.add1644.

- [209] M J P Gingras and P A McClarty. “Quantum spin ice: a search for gapless quantum spin liquids in pyrochlore magnets”. In: *Reports on Progress in Physics* 77.5 (10 May 2014), p. 056501. DOI: 10.1088/0034-4885/77/5/056501.
- [210] Owen Benton, Olga Sikora, and Nic Shannon. “Seeing the light: Experimental signatures of emergent electromagnetism in a quantum spin ice”. In: *Phys. Rev. B* 86 (7 Aug. 2012), p. 075154. DOI: 10.1103/PhysRevB.86.075154.
- [211] Gang Chen. “Spectral periodicity of the spinon continuum in quantum spin ice”. In: *Phys. Rev. B* 96 (8 Aug. 2017), p. 085136. DOI: 10.1103/PhysRevB.96.085136.
- [212] Argha Banerjee, Sergei V. Isakov, et al. “Unusual Liquid State of Hard-Core Bosons on the Pyrochlore Lattice”. In: *Phys. Rev. Lett.* 100 (4 Jan. 2008), p. 047208. DOI: 10.1103/PhysRevLett.100.047208.
- [213] Nic Shannon, Olga Sikora, et al. “Quantum Ice: A Quantum Monte Carlo Study”. In: *Phys. Rev. Lett.* 108 (6 Feb. 2012), p. 067204. DOI: 10.1103/PhysRevLett.108.067204.
- [214] Yasuyuki Kato and Shigeki Onoda. “Numerical Evidence of Quantum Melting of Spin Ice: Quantum-to-Classical Crossover”. In: *Phys. Rev. Lett.* 115 (7 Aug. 2015), p. 077202. DOI: 10.1103/PhysRevLett.115.077202.
- [215] Chun-Jiong Huang, Changle Liu, et al. “Extended Coulomb liquid of paired hardcore boson model on a pyrochlore lattice”. In: *Phys. Rev. Res.* 2 (4 Oct. 2020), p. 042022. DOI: 10.1103/PhysRevResearch.2.042022.
- [216] Yi-Ping Huang, Gang Chen, and Michael Hermele. “Quantum Spin Ices and Topological Phases from Dipolar-Octupolar Doublets on the Pyrochlore Lattice”. In: *Phys. Rev. Lett.* 112 (16 Apr. 2014), p. 167203. DOI: 10.1103/PhysRevLett.112.167203.
- [217] Jeffrey G. Rau and Michel J.P. Gingras. “Frustrated Quantum Rare-Earth Pyrochlores”. In: *Annual Review of Condensed Matter Physics* 10.1 (10 May 2019), pp. 357–386. DOI: 10.1146/annurev-conmatphys-022317-110520.
- [218] Romain Sibille, Elsa Lhotel, et al. “Candidate Quantum Spin Liquid in the Ce^{3+} Pyrochlore Stannate $\text{Ce}_2\text{Sn}_2\text{O}_7$ ”. In: *Phys. Rev. Lett.* 115 (9 Aug. 2015), p. 097202. DOI: 10.1103/PhysRevLett.115.097202.
- [219] Romain Sibille, Nicolas Gauthier, et al. “A quantum liquid of magnetic octupoles on the pyrochlore lattice”. In: *Nature Physics* 16.5 (10 May 2020), pp. 546–552. ISSN: 1745-2481. DOI: 10.1038/s41567-020-0827-7.
- [220] J. Gaudet, E. M. Smith, et al. “Quantum Spin Ice Dynamics in the Dipole-Octupole Pyrochlore Magnet $\text{Ce}_2\text{Zr}_2\text{O}_7$ ”. In: *Phys. Rev. Lett.* 122 (18 May 2019), p. 187201. DOI: 10.1103/PhysRevLett.122.187201.
- [221] Bin Gao, Tong Chen, et al. “Experimental signatures of a three-dimensional quantum spin liquid in effective spin-1/2 $\text{Ce}_2\text{Zr}_2\text{O}_7$ pyrochlore”. In: *Nature Physics* 15.10 (10 Oct. 2019), pp. 1052–1057. ISSN: 1745-2481. DOI: 10.1038/s41567-019-0577-6.
- [222] E. M. Smith, O. Benton, et al. “Case for a $U(1)_\pi$ Quantum Spin Liquid Ground State in the Dipole-Octupole Pyrochlore $\text{Ce}_2\text{Zr}_2\text{O}_7$ ”. In: *Phys. Rev. X* 12 (2 Apr. 2022), p. 021015. DOI: 10.1103/PhysRevX.12.021015.

- [223] D. R. Yahne, B. Placke, et al. *Dipolar spin ice regime proximate to an all-in-all-out Néel ground state in the dipolar-octupolar pyrochlore $Ce_2Sn_2O_7$* . arXiv:2211.15140. May 2022. DOI: 10.48550/arXiv.2211.15140.
- [224] Victor Porée, Elsa Lhotel, et al. “Crystal-field states and defect levels in candidate quantum spin ice $Ce_2Hf_2O_7$ ”. In: *Phys. Rev. Mater.* 6 (4 Apr. 2022), p. 044406. DOI: 10.1103/PhysRevMaterials.6.044406.
- [225] Victor Porée, Anish Bhardwaj, et al. *Dipolar-octupolar correlations and hierarchy of exchange interactions in $Ce_2Hf_2O_7$* . arXiv:2305.08261. May 2023. DOI: 10.48550/arXiv.2305.08261.
- [226] Anish Bhardwaj, Shu Zhang, et al. “Sleuthing out exotic quantum spin liquidity in the pyrochlore magnet $Ce_2Zr_2O_7$ ”. In: *npj Quantum Materials* 7.1 (10 May 2022), p. 51. ISSN: 2397-4648. DOI: 10.1038/s41535-022-00458-2.
- [227] Adarsh S. Patri, Masashi Hosoi, and Yong Baek Kim. “Distinguishing dipolar and octupolar quantum spin ices using contrasting magnetostriction signatures”. In: *Phys. Rev. Res.* 2 (2 June 2020), p. 023253. DOI: 10.1103/PhysRevResearch.2.023253.
- [228] Owen Benton. “Ground-state phase diagram of dipolar-octupolar pyrochlores”. In: *Phys. Rev. B* 102 (10 Sept. 2020), p. 104408. DOI: 10.1103/PhysRevB.102.104408.
- [229] Imre Hagymási, Robin Schäfer, et al. “Possible Inversion Symmetry Breaking in the $S = 1/2$ Pyrochlore Heisenberg Magnet”. In: *Phys. Rev. Lett.* 126 (11 Mar. 2021), p. 117204. DOI: 10.1103/PhysRevLett.126.117204.
- [230] J. Xu, Owen Benton, et al. “Anisotropic exchange Hamiltonian, magnetic phase diagram, and domain inversion of $Nd_2Zr_2O_7$ ”. In: *Phys. Rev. B* 99 (14 Apr. 2019), p. 144420. DOI: 10.1103/PhysRevB.99.144420.
- [231] B. Canals and C. Lacroix. “Pyrochlore Antiferromagnet: A Three-Dimensional Quantum Spin Liquid”. In: *Physical Review Letters* 80.13 (10 Mar. 1998), pp. 2933–2936. DOI: 10.1103/PhysRevLett.80.2933.
- [232] Benjamin Canals and Claudine Lacroix. “Quantum spin liquid: The Heisenberg antiferromagnet on the three-dimensional pyrochlore lattice”. In: *Phys. Rev. B* 61 (2 Jan. 2000), pp. 1149–1159. DOI: 10.1103/PhysRevB.61.1149.
- [233] Jung Hoon Kim and Jung Hoon Han. “Chiral spin states in the pyrochlore Heisenberg magnet: Fermionic mean-field theory and variational Monte Carlo calculations”. In: *Phys. Rev. B* 78 (18 Nov. 2008), p. 180410. DOI: 10.1103/PhysRevB.78.180410.
- [234] Cenke Xu. “Monopole flux state on the pyrochlore lattice”. In: *Phys. Rev. D* 79 (10 May 2010), p. 104033. DOI: 10.1103/PhysRevB.79.144432.
- [235] Sung-Bin Lee, Shigeki Onoda, and Leon Balents. “Generic quantum spin ice”. In: *Phys. Rev. B* 86 (10 Sept. 2012), p. 104412. DOI: 10.1103/PhysRevB.86.104412.
- [236] Yasir Iqbal, Tobias Müller, et al. “Quantum and Classical Phases of the Pyrochlore Heisenberg Model with Competing Interactions”. In: *Physical Review X* 9.1 (10 Jan. 2019), p. 011005. DOI: 10.1103/PhysRevX.9.011005.

- [237] Hirokazu Tsunetsugu. “Antiferromagnetic Quantum Spins on the Pyrochlore Lattice”. In: *Journal of the Physical Society of Japan* 70.3 (10 Mar. 2001), 640–643. ISSN: 1347-4073. DOI: 10.1143/jpsj.70.640.
- [238] Cenke Xu. “Valence-Bond Crystal and Anisotropic Excitation Spectrum on 3-Dimensionally Frustrated Pyrochlore”. In: *Phys. Rev. D* 67 (10 May 2010), p. 104033. ISSN: 0031-9015. DOI: 10.1143/JPSJ.67.4022.
- [239] A. B. Harris, A. J. Berlinsky, and C. Bruder. “Ordering by quantum fluctuations in a strongly frustrated Heisenberg antiferromagnet”. In: *Journal of Applied Physics* 69.8 (10 Aug. 1991), p. 5200. ISSN: 0021-8979. DOI: 10.1063/1.348098.
- [240] Erez Berg, Ehud Altman, and Assa Auerbach. “Singlet Excitations in Pyrochlore: A Study of Quantum Frustration”. In: *Physical Review Letters* 90.14 (10 Apr. 2003), p. 147204. ISSN: 0031-9015. DOI: 10.1103/PhysRevLett.90.147204.
- [241] Jürgen Schnack, Jörg Schulenburg, and Johannes Richter. “Magnetism of the $N = 42$ kagome lattice antiferromagnet”. In: *Phys. Rev. B* 98 (9 Sept. 2018), p. 094423. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.98.094423.
- [242] Imre Hagymási, Robin Schäfer, et al. “Magnetization process and ordering of the $S = \frac{1}{2}$ pyrochlore Heisenberg antiferromagnet in a magnetic field”. In: *Phys. Rev. B* 106 (6 Aug. 2022), p. L060411. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.106.L060411.
- [243] Nikita Astrakhantsev, Tom Westerhout, et al. “Broken-Symmetry Ground States of the Heisenberg Model on the Pyrochlore Lattice”. In: *Phys. Rev. X* 11 (4 Oct. 2021), p. 041021. ISSN: 0031-9015. DOI: 10.1103/PhysRevX.11.041021.
- [244] Max Hering, Vincent Noculak, et al. “Dimerization tendencies of the pyrochlore Heisenberg antiferromagnet: A functional renormalization group perspective”. In: *Phys. Rev. B* 105 (5 Feb. 2022), p. 054426. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.105.054426.
- [245] Robin Schäfer, Benedikt Placke, et al. *Abundance of hard-hexagon crystals in the quantum pyrochlore antiferromagnet*. arXiv:2210.07235. May 2022. DOI: 10.48550/arXiv.2210.07235.
- [246] Owen Benton, L. D. C. Jaubert, et al. “Quantum Spin Ice with Frustrated Transverse Exchange: From a π -Flux Phase to a Nematic Quantum Spin Liquid”. In: *Phys. Rev. Lett.* 121 (6 Aug. 2018), p. 067201. ISSN: 0031-9015. DOI: 10.1103/PhysRevLett.121.067201.
- [247] Yao-Dong Li and Gang Chen. “Symmetry enriched U(1) topological orders for dipole-octupole doublets on a pyrochlore lattice”. In: *Phys. Rev. B* 95 (4 Jan. 2017), p. 041106. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.95.041106.
- [248] Benedikt Placke, Roderich Moessner, and Owen Benton. “Hierarchy of energy scales and field-tunable order by disorder in dipolar-octupolar pyrochlores”. In: *Phys. Rev. B* 102 (24 Dec. 2020), p. 245102. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.102.245102.
- [249] Stephen W. Lovesey. *Theory of Neutron Scattering from Condensed Matter*. Vol. 2. Clarendon Press, May 1984, p. 104033. DOI: 10.1143/JPSJ.67.4022.

- [250] S. W. Lovesey and G. van der Laan. “Magnetic multipoles and correlation shortage in the pyrochlore cerium stannate $\text{Ce}_2\text{Sn}_2\text{O}_7$ ”. In: *Phys. Rev. B* 101 (14 Apr. 2020), p. 144419. ISSN: 0031-9015. DOI: 10.1143/JPSJ.67.4022.
- [251] M. Rotter and A. T. Boothroyd. “Going beyond the dipole approximation to improve the refinement of magnetic structures by neutron diffraction”. In: *Phys. Rev. B* 79 (14 Apr. 2009), p. 140405. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.79.140405.
- [252] Baoming Tang, Ehsan Khatami, and Marcos Rigol. “A short introduction to numerical linked-cluster expansions”. In: *Computer Physics Communications* 184.3 (10 May 2013), pp. 557–564. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2012.10.008>.
- [253] Robin Schäfer, Imre Hagymási, et al. “Pyrochlore $S = \frac{1}{2}$ Heisenberg antiferromagnet at finite temperature”. In: *Phys. Rev. B* 102 (5 Aug. 2020), p. 054408. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.102.054408.
- [254] Shu Zhang, Hitesh J. Changlani, et al. “Dynamical Structure Factor of the Three-Dimensional Quantum Spin Liquid Candidate $\text{NaCaNi}_2\text{F}_7$ ”. In: *Phys. Rev. Lett.* 122 (16 Apr. 2019), p. 167203. ISSN: 0031-9015. DOI: 10.1103/PhysRevLett.122.167203.
- [255] Matthew Powell, Liurukara D. Sanjeeva, et al. “Hydrothermal Crystal Growth of Rare Earth Tin Cubic Pyrochlores, $\text{RE}_2\text{Sn}_2\text{O}_7$ (RE = La-Lu): Site Ordered, Low Defect Single Crystals”. In: *Cryst. Growth Des.* 7 (10 May 2019), pp. 4920–4926. ISSN: 0031-9015. DOI: 10.1143/JPSJ.67.4022.
- [256] Anders W. Sandvik. “Stochastic series expansion method with operator-loop update”. In: *Phys. Rev. B* 59 (22 June 1999), R14157–R14160. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.59.R14157.
- [257] Roger G Melko. “Simulations of quantum XXZ models on two-dimensional frustrated lattices”. In: *Journal of Physics: Condensed Matter* 19.14 (10 Mar. 2007), p. 145203. ISSN: 0031-9015. DOI: 10.1088/0953-8984/19/14/145203.
- [258] E. M. Smith, J. Dudemaine, et al. *Quantum Spin Ice Response to a Magnetic Field in the Dipole-Octupole Pyrochlore $\text{Ce}_2\text{Zr}_2\text{O}_7$* . arXiv:2306.13183. May 2010. DOI: 10.48550/arXiv.2306.13183.
- [259] Xu-Ping Yao, Yao-Dong Li, and Gang Chen. “Pyrochlore $U(1)$ spin liquid of mixed-symmetry enrichments in magnetic fields”. In: *Phys. Rev. Research* 2 (1 Mar. 2020), p. 013334. ISSN: 0031-9015. DOI: 10.1103/PhysRevResearch.2.013334.
- [260] Yu I Dublennykh. “Ground states of the Ising model on an anisotropic triangular lattice: stripes and zigzags”. In: *Journal of Physics: Condensed Matter* 25.40 (10 Sept. 2013), p. 406003. ISSN: 0031-9015. DOI: 10.1088/0953-8984/25/40/406003.
- [261] M W Long. “Effects that can stabilise multiple spin-density waves”. In: *Journal of Physics: Condensed Matter* 1.17 (10 May 1989), pp. 2857–2874. ISSN: 0031-9015. DOI: 10.1088/0953-8984/1/17/008.
- [262] M. T. Heinilä and A. S. Oja. “Selection of the ground state in type-I fcc antiferromagnets in an external magnetic field”. In: *Phys. Rev. B* 48 (10 Sept. 1993), pp. 7227–7237. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.48.7227.

- [263] A. L. Chernyshev and M. E. Zhitomirsky. “Quantum Selection of Order in an XXZ Antiferromagnet on a Kagome Lattice”. In: *Phys. Rev. Lett.* 113 (23 Dec. 2014), p. 237202. ISSN: 0031-9015. DOI: 10.1103/PhysRevLett.113.237202.
- [264] Ioannis Rousochatzakis and Natalia B. Perkins. “Classical Spin Liquid Instability Driven By Off-Diagonal Exchange in Strong Spin-Orbit Magnets”. In: *Phys. Rev. Lett.* 118 (14 Apr. 2017), p. 147204. ISSN: 0031-9015. DOI: 10.1103/PhysRevLett.118.147204.
- [265] J. Xu, Owen Benton, et al. “Anisotropic exchange Hamiltonian, magnetic phase diagram, and domain inversion of $\text{Nd}_2\text{Zr}_2\text{O}_7$ ”. In: *Phys. Rev. B* 99 (14 Apr. 2019), p. 144420. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.99.144420.
- [266] J. Xu, A. T. M. N. Islam, et al. “Field-induced quantum spin- $\frac{1}{2}$ chains and disorder in $\text{Nd}_2\text{Zr}_2\text{O}_7$ ”. In: *Phys. Rev. B* 98 (6 Aug. 2018), 060408(R). ISSN: 0031-9015. DOI: 10.1103/PhysRevB.98.060408.
- [267] Elliott Lieb, Theodore Schultz, and Daniel Mattis. “Two soluble models of an antiferromagnetic chain”. In: *Annals of Physics* 16.3 (10 May 1961), pp. 407–466. ISSN: 0003-4916. DOI: [https://doi.org/10.1016/0003-4916\(61\)90115-4](https://doi.org/10.1016/0003-4916(61)90115-4).
- [268] T. Fennell, O. A. Petrenko, et al. “Field-induced partial order in the spin ice dysprosium titanate”. In: *Applied Physics A* 74.1 (10 Dec. 2002), s889–s891. ISSN: 1432-0630. DOI: 10.1007/s003390201638.
- [269] T. Fennell, O. A. Petrenko, et al. “Neutron scattering studies of the spin ices $\text{Ho}_2\text{Ti}_2\text{O}_7$ and $\text{Dy}_2\text{Ti}_2\text{O}_7$ in applied magnetic field”. In: *Phys. Rev. B* 72 (22 Dec. 2005), p. 224411. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.72.224411.
- [270] J. P. Clancy, J. P. C. Ruff, et al. “Revisiting static and dynamic spin-ice correlations in $\text{Ho}_2\text{Ti}_2\text{O}_7$ with neutron scattering”. In: *Phys. Rev. B* 79 (1 Jan. 2009), p. 014408. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.79.014408.
- [271] O. Derzhko and T. Krokhamlskii. “Numerical Approach for the Study of the Spin- $\frac{1}{2}$ XY Chains Dynamic Properties”. In: *physica status solidi (b)* 208.1 (10 May 1998), pp. 221–248. ISSN: 0031-9015. DOI: [https://doi.org/10.1002/\(SICI\)1521-3951\(199807\)208:1<221::AID-PSSB221>3.0.CO;2-E](https://doi.org/10.1002/(SICI)1521-3951(199807)208:1<221::AID-PSSB221>3.0.CO;2-E).
- [272] Rajiv R. P. Singh and Sudip Chakravarty. “High-temperature series expansion for spin glasses. I. Derivation of the series”. In: *Phys. Rev. B* 36 (1 July 1987), pp. 546–558. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.36.546.
- [273] Rajiv R. P. Singh and Sudip Chakravarty. “High-temperature series expansion for spin glasses. II. Analysis of the series”. In: *Phys. Rev. B* 36 (1 July 1987), pp. 559–566. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.36.559.
- [274] Walter Marshall and S W Lovesey. *Theory of thermal neutron scattering*. en. Vol. 67. Monographs on Physics. Melbourne, VIC, Australia: Titles Distributed by Oxford University Press, Sept. 1971, p. 104033. DOI: 10.1143/JPSJ.67.4022.
- [275] L. D. Landau und J. M. Lifschitz. “Theory of the dispersion of magnetic permeability in ferromagnetic bodies”. In: *Phys. Z. Sowjetunion*. 8 (10 May 1935), p. 153. ISSN: 0031-9015. DOI: 10.1143/JPSJ.67.4022.

-
- [276] P. H. Conlon and J. T. Chalker. “Spin Dynamics in Pyrochlore Heisenberg Antiferromagnets”. In: *Phys. Rev. Lett.* 102 (23 June 2009), p. 237206. ISSN: 0031-9015. DOI: 10.1103/PhysRevLett.102.237206.
- [277] Mathieu Taillefumier, Julien Robert, et al. “Semiclassical spin dynamics of the antiferromagnetic Heisenberg model on the kagome lattice”. In: *Phys. Rev. B* 90 (6 Aug. 2014), p. 064419. ISSN: 0031-9015. DOI: 10.1103/PhysRevB.90.064419.
- [278] D. Loison, C. L. Qin, et al. “Canonical local algorithms for spin systems: heat bath and Hastings’s methods”. In: *The European Physical Journal B - Condensed Matter and Complex Systems* 41.3 (10 Oct. 2004), pp. 395–412. ISSN: 1434-6036. DOI: 10.1140/epjb/e2004-00332-5.
- [279] Peter J Brockwell and Richard A Davis. *Introduction to time series and forecasting*. en. 2nd ed. Vol. 67. Springer Texts in Statistics. New York, NY: Springer, Dec. 2002, p. 104033. DOI: 10.1143/JPSJ.67.4022.
- [280] R Kubo. “The fluctuation-dissipation theorem”. In: *Rep. Prog. Phys* 29.1 (10 Jan. 1966), pp. 255–284. ISSN: 0031-9015. DOI: 10.1088/0034-4885/29/1/306.

Versicherung

Hiermit versichere ich, dass ich die vorliegende Arbeit ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe; die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sind als solche kenntlich gemacht. Die Arbeit wurde bisher weder im Inland noch im Ausland in gleicher oder ähnlicher Form einer anderen Prüfungsbehörde vorgelegt.

Die Arbeit wurde am Max Planck Institut für Physik Komplexer Systeme in Dresden angefertigt. Wissenschaftlicher Betreuer war Roderich Moessner.

Benedikt Andreas Placke
August 28, 2023, Dresden